[Telairity](#) has made a name for itself as the industry's leading video processing solutions provider.
Please write to us at sales@telairity.com to learn more about our products and to collaborate with our team.

A Deep Dive Into
UHD Technology

A TELAIRITY WHITE PAPER

Harlan McGhan
TELAIRITY | 3375 SCOTT BLVD., SUITE 342, SANTA CLARA, CA 95054

Table of Contents

A Deep Dive Into UHD Technology

Just when we might think that so-called “Full HD” resolution (or “2K”) was the absolute cutting edge, “Ultra-High Definition” (UHD) made an appearance and changed the equation. In fact, 4K UHD technology has been in the news since 2010, but 2015 saw a steep drop in the prices of devices supporting 4K and, judging by the latest sales figures, the market has warmed considerably to the new resolution standard. But, just as 4K is starting to replace HD, some broadcasters are beginning to promote 8K as the natural and inevitable successor to 4K—an upgrade that apparently cannot arrive too soon, given its advantages over 4K.

What in the world is going on? The SD standard lasted over half a century. Now, a new resolution standard (HD, 4K, 8K) seems to be appearing every decade or so. Is 8K to be followed by 16K, which in turn will be superseded by 32K? Is there any end to this relentless upgrade cycle of digital resolution standards?

1 What is Digital UHD Technology

1.1 First, Resolution Is About Pixel Count

In the old analog world of CRT display screens, screens were measured by the number of scan lines they supported. Before HD, every TV had either 480 or 576 visible scan lines (depending on the SD standard used, NTSC or PAL). The switch to high-definition precipitated a shift in the underlying technology, from analog to digital.

In the new digital world of TV technology, display screens are no longer measured by the number of scan lines they support, but rather by the number of pixels they display. Digitally speaking, the highest resolution SD screens are now pixel arrays 720 wide x 480 or 576 high. Full HD screens are pixel arrays 1920 wide x 1080 high. In round numbers, HD resolution is about 2000 x 1000, which gets shortened to “2K” in digital-speak. 4K UHD simply doubles each HD (2K) dimension, to about 4000 x 2000 (in exact terms, to 3840 x 2160). 8K UHD is a further doubling of the pixel array over 4K, to about 8000 x 4000 (in exact terms, to 7680 x 4320). Roundly speaking, then, HD is about 2 million pixels/screen, while 4K is about 8 million pixels/screen, or 4X the resolution of HD, and 8K is about 32 million pixels/screen, or 4X the resolution of 4K (16X the resolution of HD).

The figure below graphically compares resolution/size and pixel counts per frame for all resolution standards from SD to 8K UHD.

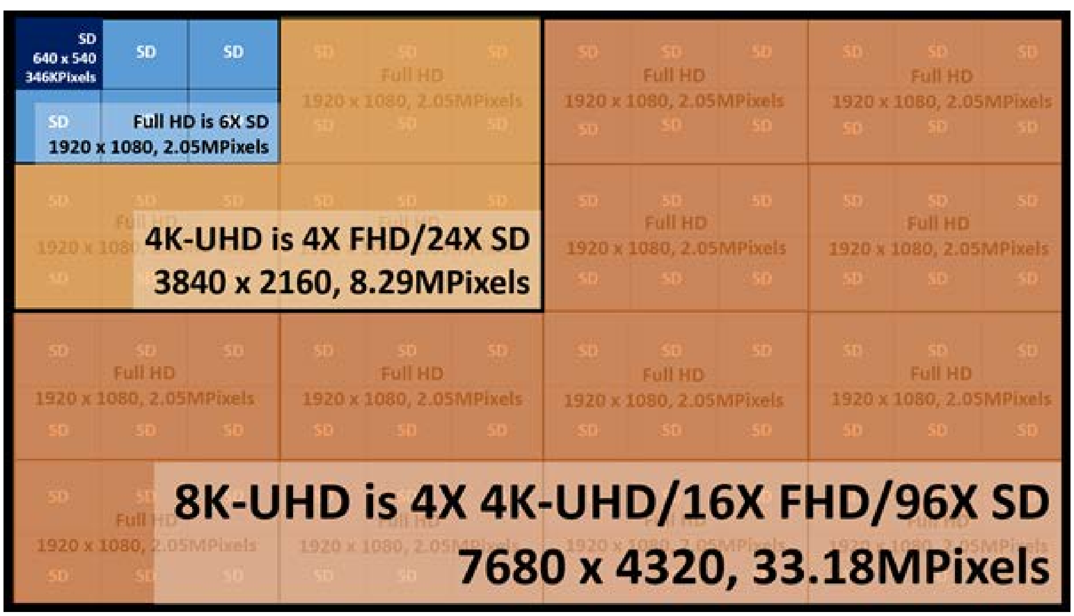

1.2 Much Remains Outside Pixel Count

It is easy to get confused here because pixel count is only one aspect of the technology used to manufacture displays. Equally critical is the technology used to render pixels (whatever their number). Rendering technology introduces words like “Plasma” and acronyms like “LCD” (Liquid Crystal Display), “LED” (Light Emitting Diode), “OLED” (Organic LED), “QLED” (Quantum Dot LED), and so on. It also controls the maximum darkness and lightness of a screen (its black/white contrast ratio), as well as how bright and vivid colors appear. Still another issue has to do with screen shape (curved or flat), and the effect this has on the viewing experience.

Obviously, if you change multiple aspects of a display at once, the total impact on the viewing experience can be far greater than the impression made by any one change in isolation. No doubt, shifting from a flat LCD HD display to a curved OLED UHD display will dramatically transform one’s viewing experience. But what part of this transformation does the change in pixel count (i.e., the shift from HD to UHD) contribute by itself, distinct from the other independent new technologies for displays (like rendering technology and screen shape) now coming into commercial use?

1.3 Second, Resolution Specifies Data Rates

The important point about digital bitmap formats like HD and UHD is that they fix the number of pixels displayed in each video frame, independent of screen size. Every “full HD” screen is an array of 1920 x 1080 pixels, whether the screen measures 30” or 70” or some other number. Similarly, every 4K UHD screen is an array of 3840 x 2160 pixels, regardless of how large or small the display.

The standard also specifies the format of individual pixels in the array, i.e., whether the pixels are “8 bit”, “10 bit”, or some other number of bits. Multiplying these two numbers together—pixels-per-frame x bits-per-pixel—will generate a bits-per-frame number. For example, using “8-bit” pixels, a full HD frame is 1920 x 1080 = 2,073,600 pixels x 24 bits/pixel[1] = 49,766,400 bits.

The last piece of the data rate puzzle is for the standard to specify a frame per second or fps value for video. The long-time motion picture standard is 24 fps, but both the NTSC SD and

---

[1]Why are pixels 24 bits long described as “8-bit”? It’s because “8-bit” refers not to pixel length, but rather to “channel” length, or the number of bits used to encode each of the 3 primary colors (Red-Green-Blue) that make up a pixel. Adding the 3 8-bit primary color “channels” together gives the overall total of 3 x 8 or 24 bits/pixel.

Note that while the pixel array (W x H) and the length of a pixel is fixed by the resolution standard, the standard does not fix pixel size/shape. These “local” pixel attributes are determined by the display screen receiving the data, and will vary widely from small “retina” displays to big screen TVs. Even specific color specifications (that is to say, the particular 8+8+8 RGB data for individual pixels), will render differently on different types of displays.

1080i “full HD” standards use 30 fps, while the PAL SD standard uses 25 fps. The 720p HD standard uses 60 fps, the same as the 4K UHD standard.

Continuing with our HD example, to get the total bits per second (bps) data rate, we have to multiply the roughly 50 million bits/frame of full HD by 30 fps, which yields 1.5 billion bps in round numbers (or precisely: 1,493,292,000 bps).

1.4 Pixel Size and PPI Fixed by the Display

Given a fixed number of pixels—roughly, 2M in a 2K x 1K “2K” HD bitmap, 8M in a 4K x 2K “4K” UHD bitmap—what must happen as an HD or UHD screen gets larger or smaller is that the individual pixels in the array must grow or shrink accordingly. This fact brings us to yet another critical metric for displays, known as ppi or pixels-per-inch. Although an old idea (familiar to anyone who has ever bought a raster printer as dpi or dots per inch), this metric was first popularized for displays by Apple, with the term “retina display.” This is a display in which the pixels are too small to be individually distinguished by the human eye, even on relatively close handheld viewing. In ppi terms, pixels get too small to be seen somewhere just short of the number 300, so a “retina display” is any screen with a ppi number of 300 or greater. Note that 300 is also the classic dpi number for quality printing/scanning.

1.5 “Recommended Viewing Distance”

The notion of ppi, in turn, brings us to our last critical metric, viewing distance. Even the largest pixels can be made too small to be individually distinguished by the human eye, by the simple expedient of moving the eye further away from the display. Distance is the principle behind “Jumbotron” displays, which have pixels the size of playing cards (or bigger) but are designed to be viewed from hundreds of feet away.

Applying a retina display standard to TV screens would make them disappointingly small. An HD screen, built to the “retina display” threshold of 300 ppi, would measure less than 7 x 4 inches (about the same as many current smartphone screens). Even a 4K UHD “retina display” would be less than 13 x 8 inches, or about the size of a tablet screen.

The reason TV screens of 50” and more are common is just that TVs are not designed for close “retina display” viewing. As screens (pixels) are made bigger, the adjustment made by display manufacturers is simply to increase the recommended viewing distance (thereby maintaining a constant apparent pixel size in the eye of the viewer). Conversely, as screens (pixels) get smaller viewers are allowed to move closer, again following recommended viewing distance guidelines, with no change in the apparent pixel size.

1.6 Basic Difference Between HD and UHD

In a nutshell, then—ignoring, for the moment, differences in pixel length ("8-bit" vs. "10-bit") and frame rate (30 vs. 60 fps)—the whole technical difference between an HD display and a UHD display reduces to relative pixel/screen size. Since 4K UHD formats cram 4X the number of pixels onto a screen as HD, for screens of the same size, 4K pixels are ¼ the size of HD pixels; conversely, for pixels of the same size, 4K screens have 4X the viewing area of HD screens. The same relationships hold between 4K and 8K UHD formats. If screen size stays constant, 8K pixels have to shrink to ¼ the size of 4K pixels. If pixel size remains constant, 8K screens have to expand to 4X the area of 4K displays.

Thus, in basic terms (keeping pixel lengths and frame rates constant), the whole viewing difference between an HD display and a UHD display comes down to just one point: bigger screens with no loss of visual quality—where "visual quality" is measured by the single metric of apparent pixel size. It makes no difference whether you replace your old HD TV with a new UHD TV of the same screen size, and move 4X closer to it, or keep the same viewing distance, but replace your old TV with a 4X bigger UHD display. In both cases, the effect is the same: the screen looms 4X larger in your visual field.

1.7 A More Immersive Viewing Experience

The ability to increase apparent screen size with no loss of visual quality is not everything, but it is not nothing, either. The apparent size of a screen in our viewing area is a key factor in what is called viewing *immersion*; indeed, covering one's entire visual field with a screen is the principle responsible for creating the illusions of virtual reality.

On this analysis, then, the advantage of UHD over HD is primarily its ability to create a more immersive viewing experience, by allowing the viewer to get closer to screens of the same size, and view larger screens at the same distance, with no loss in visual quality. This increase in apparent screen size is presumably a good thing, at least when we want to be more immersed in what we are viewing. But, like many good things, UHD has its trade-offs.

1.8 The Problem with Digital Video

The most obvious trade-off for 4K UHD is simply the cost of quadrupling the number of pixels per video frame, from about 2 million to about 8 million. As a viewer, you might think that doesn't matter, as long as advancing display technology makes new 8-million pixel 4K screens available in the same price range formerly paid for comparable 2-million pixel HD screens. Like an iceberg, however, the implications of multiplying pixels run far deeper than the visible surface of a 4K screen.

Digitally speaking, every pixel is a number, specifically a binary number that represents a particular color shade. For each pixel, the display reads its number and generates the colored block appropriate for that number, in the location appropriate for that pixel, in a size appropriate to the screen's dimensions and resolution format.

The pixel numbering standard in common use today for broadcast television is so-called "8-bit" color, which generates a 24-bit 3-channel RGB binary number for each pixel, sufficient to enable a palette of over 16 million colors. Since 16 million is more RGB color shades than even the most discerning human eye can distinguish, 8-bit color (24 bits/pixel) is sometimes called "true" color, as the first and simplest digital color scheme to enable everything the human eye can see (and more).[2]

The problem created by digital imagery in general, and HD and UHD television in particular, isn't that digital technology is inferior to older analog technology, or that it is inadequate to express the full range of our senses. The issue is just that digital technology able to provide a high-quality experience takes a lot of bits, and improvements in quality take even more bits.

Specifically, an HD picture composed of 2 million pixels, each corresponding to a 24-bit number, requires 48 million bits to express.[3] And that is just for a single frame. Full HD plays out at 30 frames a second, meaning a total bit rate of nearly 1.5 billion bits every second. 1.5 billion bits per second (1.5Bbps) is not just a large number; it is an overwhelming number. As a general rule, it is impractical to store 1.5 billion bits for every second of HD video captured, let alone transmit bits at that rate.

2 Importance of Video Compression

Fortunately, there is a powerful remedy for the proliferation of bits required by digital rendering technology, namely digital compression technology. Compression technology is especially powerful for video, where standards like H.264 allow the elimination of 299 bits out of every 300, reducing 1.5 billion bits a second to a much more manageable 5 million bits a second (or even lower, in some cases). But what happens to data rates when the television industry shifts from HD to higher (then still higher) UHD resolution standards?

[2]For further discussion of digital color coding, see Appendix 1: Binary Color Coding.

[3] 48 million is the "round number" approximation of 2 million times 24. Multiplying out the actual HD numbers (2,073,600 pixels/frame x 24 bits/pixel) yields an actual answer much closer to 50 million (49,766,400).

2.1 How Data Compression Works

Broadly speaking, digital data compression is simply a process of eliminating information in inverse order of importance. More precisely, digital compression reduces to two logical tasks. The first and most straightforward task is simply to remove everything redundant in the data, i.e., to find the smallest number of bits that can be used to encode any given amount of data with no loss of information. Technically, this is called *entropy* encoding.

The second task begins with a ranking exercise, sorting the non-redundant or meaningful data in order of its interest or importance. Data is then eliminated starting at the bottom, with the least significant (most uninteresting), and working up through the ranks of data in order of increasing significance. This second type of compression stops after reaching a set goal, which may be either a size target or a specified level of data importance (quality standard).

The first type of compression, which retains all information and eliminates only redundancies, is known as “lossless” compression; the second type of compression, which also removes the least important or most uninteresting parts of the original data, is called “lossy” compression. For images, lossless compression alone is usually inadequate; i.e., unable to achieve the kinds of very substantial reductions needed to cope with the overwhelming number of bits generated by digital imaging. For this reason, all standard image compression technologies (JPEG, MPEG, etc.) are lossy.

The good news for video is that the first lossy reductions are imperceptible. Image sensors capture color information the human eye is not able to discriminate, so the first step in video compression is simply to eliminate what we cannot discern. If the compression process stops here, there is absolutely no degradation of visual quality. Everything the human eye can see is still present in the remaining data.

If still more compression is needed, the next step is to eliminate what perfect vision can see, but whose absence most people will not notice. From there, the process proceeds to losses more widely seen, but usually regarded as insignificant, and so on.

Indeed, the rest of the story of declining bitrates should be familiar to anyone who has watched much streaming video. As the discarded information becomes more and more critical, the picture begins to soften up. Eventually, the edges of objects blur and run together. Finally, the objects themselves become unrecognizable, dissolving first into a mosaic of colored blocks, then into nothing at all.

2.2 The Limits of Video Compression

The point to this story of progressively degrading image quality with increasing levels of compression is just that there are limits to what even the best possible compression technology can achieve without significant degradation. Great video compression technology, like H.264, may be able to eliminate 299 out of every 300 bits in the original image, with little or no perceptible loss in visual quality. Even better compression technology might be able to eliminate still more bits, leaving only 1 in 400 or 500. But no compression technology can remove all the bits, and no compression technology can delete bits from the hard-residual core of important information (however large or small it may be) without significantly degrading the quality of what remains.

There is an important moral here about what to expect from the next step in video compression technology— the move from H.264 (MPEG-4 or AVC, for Advanced Video Coding) to H.265 (MPEG-5 or HEVC, for High-Efficiency Video Coding). We will return to this issue later.

3 UHD Data Rates

3.1 The Basic Calculation

For now, however, the point is simple. 4K UHD has 4X the pixel count of HD, and pixels are nothing but 24-bit bundles of color data. So the starting point for any discussion of upgrading HD to UHD *without* degrading picture quality (cutting more of the *significant* information from the bit stream) is how to move four times as many bits across our broadcast networks. More precisely, the uncompressed HD figure of about 1.5 billion bits per second (1.5Gbps) jumps to about 6 billion bits per second (6Gbps) for 4K. Similarly, the compressed HD figure of about 5 million bits per second (5Mbps)—barring any substantial improvements in compression technology—jumps to about 20 million bits per second (20Mbps). [4]

In truth, shifting from HD to 4K UHD at the cost of moving just 4X the amount of information would be the good news. The bad news is that 4K UHD requires far more than just a mere quadrupling of data rates.

[4] In brief, image compression is very much like floating a boat that threatens to sink in heavy seas under the weight of too much cargo. You begin by throwing overboard everything extra, and follow by throwing overboard everything dispensable, until a critical core of functionality is reached past which the boat ceases to be viable. Nothing in this process of lightening the load, however, changes the size of the boat. If you leave port aboard the UHD supertanker, you must arrive aboard the UHD supertanker, and not on some trim little yacht.

3.2 The Complete Calculation

The data challenge posed by 4K UHD is worse than a mere 4X increase in data rates because the UHD standard encompasses more than just doubling and redoubling pixel count. The UHD standard also allows an upgrade from 8-bit to 10-bit color[5] and an increase from a rate of 30 frames per second (30 fps) to a rate of 60 fps. Each of these changes further boosts the data rates required for UHD video.

The move from 8-bit to 10-bit color channels increases pixel length from 3 x 8 = 24 bits to 3 x 10 = 30 bits. Which is to say, the new UHD color standard requires a further 25% increase in bit rates. This additional 25% bump means raw data rates will not be just 6Gbps, but rather 7.5Gbps, to accommodate the six extra color bits for each pixel.

And that increase is not the end of the story, given the UHD standard also allows doubling the HD frame rate. Transmitting twice as many frames a second requires a further doubling of the bit rate, so the 7.5Gbps need to send 10-bit 4K UHD at 30 fps becomes, instead, a data rate of 15Gbps, to send the new UHD data at 60 fps. So, the cumulative effect of all three changes is that the 4X penalty in bit rates, imposed by the increased pixel counts, balloons into a 4 x 1.25 x 2 = 10X increase in total bit rates.

In short, the shift from full HD to the full 4K UHD standard requires an order of magnitude jump in data rates. At the level of uncompressed data, that's a leap from1.5Gbps to 15Gbps. At the compressed level—again, barring significant improvements in compression technology—it's a jump from 5Mbps to 50Mbps.

Is an order of magnitude change in data rates a problem? It all depends on where you fall in the video chain. For some, UHD is all gain. For others, it is a massive headache.

4 Who Benefits from UHD

4.1 Display Manufacturers Get a New Market

Clearly, display manufacturers are the group that is most enthusiastic about the new "4K" UHD standard. Given the ability to manufacture a screen with 8 million pixels, each able to display a broader range of "10-bit" HDR[6] colors at 60fps, the advantage to embracing the new UHD standard for this group is evident. Just as HD converted a mature, no-growth market for SD

[5]For further discussion of 8-bit vs. 10-bit color, see Appendix 1: Binary Color Coding.

[6]For further discussion of HDR (High Dynamic Range) color, see Appendix 1: Binary Color Coding.

screens into a high-growth market for HD screens a decade ago, 4K UHD now has the same promise to revitalize an increasingly stagnant market for HD displays.

Of course, the responsibility borne by display manufacturers for the 4K standard ends with providing a 4K UHD-capable screen. Is there any content available able to take advantage of 4X higher resolution, 2X faster frame rates, and a 1.25X increase in color data? For display makers, at least, content delivery is not a problem—or, at least, not beyond the need to help foster the conviction that UHD content is coming, as necessary to encourage screen sales. The importance of UHD to a display manufacturer is not that UHD helps create a more immersive viewing experience. Rather, it is simply that the new standard makes all HD TVs obsolete, hopefully precipitating a whole new buying cycle for UHD TVs.

There is a cautionary lesson here. Once the 2K-to-4K transition is largely over, the incentive for display manufacturers is to begin promoting whatever new standard advancing display technology will then make possible. From this viewpoint, the best news is that the new 8K UHD resolution standard is already on the books, presumably just waiting for the screen technology needed to implement arrays of 7680 x 4320 pixels, or roughly 32 million pixels.

Of course, the lurking 8K UHD standard does not stop with quadrupling the pixel count of 4K UHD, any more than 4K stopped with quadrupling the pixel count of 2K HD. It also pushes on from 10-bit to 12-bit color channels (36-bit pixels), and allows another doubling of the frame rate, to 120 fps. As a result, the increase in data rates for full 8K, over full 4K, is another 3-term multiplication problem: 4 x 1.20 x 2 = 9.6X. Piled on top of the 10X increase over HD data rates required by full 4K, then, full 8K represents a 96X increase in data rates over HD.

But even a near-100X increase in current HD data rates is not a problem for display companies since their business stops at selling the display. Figuring out how to transmit two orders of magnitude more data to new 8K screens is someone else's problem. As long as technology supports the manufacture of still higher resolution displays, and customers looking to buy a new TV set believe higher resolution always beats lower resolution, there is no reason for display manufacturers to want the drive towards a new, higher resolution standard every decade or so to ever stop. ("16K" anyone?)

4.2 OTT Providers Get a New Marketing Tool

From display manufacturers, let's turn to a group directly involved with 4K content: so-called "Over The Top" (OTT) video providers, like Netflix, Amazon, Hulu, YouTube, etc. OTT companies are in the business of acquiring video content, which they then make available over the Internet, either freely or to paid subscribers (according to their business model). Netflix has also led an OTT movement into content creation, with original series like "House of Cards."

Netflix, in particular, has been active in promoting UHD, even pledging to create all its new original content to the 4K standard. More, it has raised its own 4K ante, by producing the 3rd season of House of Cards in an upscale "6K" format (6144 x 3160 pixels). Moving to 6K production is a significant investment. Not only are 6K cameras expensive, but post-production editing costs also go up considerably with higher resolutions. The resulting source files are also huge—the 6K master copy for a 55-minute House of Cards episode is said to be a whopping 5.5 TeraBytes (TB).

Still, in the overall budget for a high-profile TV series—compared to salaries, location costs, marketing, etc.—the added production costs for 4K and higher formats are doubtlessly a relatively minor line item. And, in an era when a TB of storage costs less than $30, even the expense of storing a TB for every 10 minutes of video is trivial. OTT companies can doubtless write off the entire added cost of UHD production as a line item under marketing expenses, as part of their effort to gain a competitive advantage for their program lineup.

However, the real nub of the crises created by new UHD resolution standards is not the cost of producing, editing, or storing UHD video, but rather the challenge of delivering UHD programs, like House of Cards, in an actual 4K format. And, here, Netflix is on solid ground, because (like other OTT companies) they don't deliver video.

It may be helpful to keep this point in mind the next time you hear or read glowing comments about the imminent advent of UHD video from Netflix or other OTT providers. Certainly, they are ready to acquire and store UHD content. They may even be eager to produce their original programs in UHD formats. But their responsibility for UHD ends with embracing the new technology, for whatever marketing advantage they gain by offering more and better UHD content.

Of course, OTT companies must also provide the video they offer. Netflix does this either via mail, in a stored media format (DVD, Blu-Ray, Ultra Blu-Ray), or by streaming it over the Internet.[7] But OTT providers are not responsible for anything to do with either of these delivery ecosystems.

Thus, the Blu-ray Disc Association has a long list of contributors to the Blu-ray standard, but this list does not include Netflix. In other words, since companies interested in stored media have developed Ultra Blu-ray for 4K content, Netflix is happy to rent these disks. But, if stored media companies fail to develop a format for, say, 8K resolutions, well, that's not Netflix's call.

Much the same applies to Internet delivery. True, Netflix is easily the largest consumer of Internet bandwidth and, together with fellow OTT traveler, YouTube (a distant second in bandwidth consumption), now accounts for an astonishing over-one-half share of all Internet

[7]For further discussion of Internet streaming, see Appendix 2: Adaptive Bitrate Streaming.

traffic during prime loading periods. But OTT companies (unlike Al Gore) do not claim any credit for creating the Internet, and (like Al Gore) they take no responsibility for either maintaining or upgrading it.

If the Internet bogs down and OTT streaming stops thereby, well, your OTT provider sincerely regrets the interruption in your service, and suggests you try again later. And, if that doesn't fix the problem, you are, of course, free to call your local Internet service provider to complain.

4.3 Real-Time Data Delivery is the Real Crises

In summary, for display manufacturers, UHD is all opportunity. For forward-thinking video providers, UHD is the next step in the evolution of video technology; a step they may be not only anxious to prove they are ready to take but also eager to been seen as actively leading. The Ultra BluRay optical disk format solves the challenge of providing 4K content via stored media. Presumably, the Blu-ray Disk Association can develop an 8K format as well, should that format supersede 4K UHD. The real crisis posed by UHD only appears when we turn to the task of actually moving the massively increased amounts of data generated by UHD programs to UHD receivers in real time across remote networks.

5 How to Deliver UHD Data in Real Time

We can glimpse the height and breadth of the real-time roadblock preventing the dawn of the rosy 4K future being promoted today by TV set manufacturers and OTT content provider by returning to the figure quoted above: Netflix and YouTube together account for over half of all peak Internet traffic today. Replacing even 15% of that traffic with 10X more extensive UHD data would generate *more traffic than the entire Internet handles today*. Leaving no capacity for anyone, not on Netflix and YouTube, to do anything at all; not to mention no room for the remaining 85% of the Netflix/YouTube audience to move up to UHD.

5.1 HD to SD No Precedent for HD to UHD

If that reflection is not sufficiently sobering, consider our one and only historical precedent for a transition from HD to UHD, namely, the transition from SD to HD. Serious research into new HD formats for commercial television started back in the 1960s, with attempts to demonstrate viable HD systems beginning in the 1970s and continuing through the 1980s. The pertinent point for this discussion is that all these efforts, based on existing over-air analog technologies, ultimately failed for one and the same reason: there was no practically feasible and widely available way to move four to six times as much data for a new HD format across existing analog channels.

The development of new digital technology in the 1990s finally resolved this 30-year impasse. Critical for HD was the fact that, among other advantages, digital imaging allows the use of data compression. Of course, data compression did not, in fact, remove the bandwidth bottleneck that had stifled earlier HD initiatives; rather, it choked the data that needed to be transmitted down to a size that would fit through the new digital pipes. Indeed, data compression was so successful at reducing bandwidth needs that, during 2005-2010, the FCC reassigned bandwidth, formerly allocated to broadcasters, to Sprint/Nextel for cellular use.[8]

The bad news for the HD to UHD transition is that the analog to digital conversion is a one-trick pony. There is no similar technical legerdemain now waiting in the wings to enable an upgrade to UHD. In today's all-digital world, the only relevant difference between HD and UHD is the difference between some bits and an order of magnitude more bits.

5.2 The HD Model, Doubled and Redoubled

The optimism that underlies UHD, then, is not a belief in some fundamental technological shift, but rather the conviction that digital technology is highly elastic, able to accommodate rising numbers of bits through many doubling cycles.

Of course, there is a well-known and very practical demonstration of the ability of technology to expand in precisely this way, namely, Moore's Law, which postulates that semiconductor technology can double the number of transistors on a chip every couple of years. Over the course of the last 57 years, Moore's law has carried chip makers across 33 doubling cycles, from a single transistor in 1959 to over 8 billion transistors today, with no clear end yet in sight.

---

[8]The so-called "2 GHz relocation" of 7 Broadcast Auxiliary Service (BAS) channels. These were 17/18-MHz wide analog channels, located between 1990 and 2110 MHz. Broadcasters were allocated this radio-frequency band for internal communication between studios and outside locations (aka "backhaul" channels, to fixed transmitter sites, fixed field cameras, or mobile field reporters). The Sprint/Nextel initiative converted these 7 analog channels to 7 narrower 12-MHz wide digital channels, located in the upper 70% of the old analog band (between 2025.5 and 2109.5 MHz). The excluded frequencies, in the lower 30% of the band between 1990 and 2025, were then reassigned to Sprint/Nextel, to expand its adjacent Personal Communications Service (PCS) band (which ran between 1850-1990 MHz). As incentive for surrendering 30% of their analog BAS bandwidth, Sprint/Nextel bought new digital equipment for all the broadcasters forced to convert to the reassigned 12 MHz digital channels, to continue using auxiliary services.

The power of digital compression can be seen in the fact that the old analog BAS channels, with 17-18 MHz of bandwidth, were unable to accommodate even a single HD channel; while, using advanced MPEG-4 compression (H.264/AVC) over the narrower 12 MHz digital band, broadcasters were easily able to accommodate 2 HD channels at the same time (each at a relatively robust data rate of 6 Mbps). As a consequence, in the 2005-2010 timeframe, pioneering local stations who upgraded their broadcast studios and operation centers to HD, were able to conclude an "all HD" news upgrade. Using digital HD cameras, MPEG-4 encoders, and digital transmitters and receivers purchased with Sprint/Nextel money, they were also able to shift their field reporting (aka "ENG" for Electronic News Gathering) to the new HD standard.

For digital video, the postulated elasticity of bit transmission must be founded on some combination of the abilities to contract the rising flood of bits, through improved digital compression technology, and to increase digital channel capacity, to accommodate however many bits remain after compression.

5.3 Prospects for Better Data Compression

Clearly, the best way to deal with the rising floodtide of bits created by UHD would be by compensating improvements in data compression technology. In large part, this *is* the story of the digital SD to digital HD transition. The switch from MPEG-2 to the more aggressive set of compression tools in MPEG-4 successfully reduced the flood of additional bits created by HD to something closer to a trickle (where it did not eliminate it entirely).

To take just one data point, with MPEG-2 technology, CableLabs recommended minimum video streaming rates of 3.75Mbps for SD, but 15Mbps for HD, i.e. a 400% increase to enable the new resolution standard. However, using MPEG-4 technology, Netflix now streams full HD at a maximum rate of 5.8Mbps, a relatively modest increase of just 55% over the 3.75Mbps rate recommended for MPEG-2 SD. And even that 55% increase can be drastically cut by using an "HD Lite" format, which Netflix streams at a minimum of just 2.35Mbps. The bottom line is simply that better compression technology largely canceled the overall impact of the SD to HD transition on bitrates.

5.3.1 Historical Overview of Video Compression

There is, therefore, historical precedent for thinking that a similar scenario may play out for the HD to UHD transition. It is worth briefly pausing here to review the history of digital video compression. Although useful digital data compression dates back to the very dawn of the computer era, 1951[9], digital video compression is much newer. The Moving Pictures Expert Group (MPEG) was not formed until 1988 and did not issue its first MPEG-1 standard (aimed at CIF-format video conferencing) until 1993. MPEG-2, the standard aimed at broadcast video (digital "full SD" or 720 x 480/576 interlaced formats), appeared the following year.

Although MPEG-2 was successfully upgraded to handle new HD resolutions (both 720p and 1080i)—leading to the abandonment of a proposed MPEG-3 standard specifically for HD

---

[9]1951 is both the year the first commercial computer, UNIVAC I, was sold to the U.S. Census Bureau, and the year an MIT doctoral student named David Huffman developed the first successful algorithm for generating prefix-free variable length codes (aka "Huffman" codes). The basic idea behind Huffman's variable length coding (VLC) scheme is simple: use the shortest binary codes for the data that appears most frequently. Video compression standards have refined this idea, first as CAVLC (Context Adaptive Variable Length Coding), then as CABAC (Context Adaptive Binary Arithmetic Coding). However, Huffman's key insight remains at the heart of lossless or "entropy" compression.

broadcasting—the resulting bitrates for high-quality video tended to overwhelm transmission resources. As mentioned above, 15 Mbps was recommended as a minimum for digital cable, while the standard ATSC rate for over-the-air broadcasting was set even higher, at 19.4 Mbps. Lowering these MPEG-2 bitrates by reducing video quality tended to defeat the whole purpose of HD. Of course, this is not to say no one ever did it, as some cable providers scrambled to provide HD services by streaming MPEG-2 in the 10-14 Mbps range.

The saving grace for HD turned out to be the fact that it was slow to ramp. Although the HD standard itself dates to 1996, with the first HD broadcast occurring in 1998, HDTV penetration of U.S. households did not pass even the 10% mark for a decade (2007). This provided time for the MPEG committee to develop a robust set of more aggressive compression tools, organized under a new MPEG-4[10] standard. MPEG-4 is also known by the acronym AVC (for Advanced Video Coding), as well as by its ITU-T designation, H.264. The timely arrival of MPEG-4 compression, with its ability to cut HD data rates back to near-SD levels, prevented HD programming from ever generating a widespread bandwidth bottleneck crisis for transmission of video signals, whether beamed to satellites, sent across cables, streamed on the internet, or broadcast over the air.

Are we now poised to repeat this historical success story one more time? In fact, the Motion Pictures Experts Group released a next-generation MPEG-5 upgrade in 2013. Like MPEG-4, MPEG-5 is also known by an acronym, as HEVC (for High-Efficiency Video Coding), and by its ITU-T designation, H.265. The 10-year intervals between MPEG-2 (1994), MPEG-4 (2003, 2005), and MPEG-5 (2013) seem too precise to be just a coincidence. Presumably, just as transitioning to the new HD resolution standard was facilitated by the timely arrival of MPEG-4, so transitioning to the new 4K UHD resolution standard will be eased by the debut of MPEG-5.

Sadly, the short answer to the question, whether MPEG-5 can repeat the success of MPEG-4 in stemming the rising floodtide of bits caused by a new resolution standard, is "no."

5.3.2 Limitations of MPEG-5 (HEVC) Compression

Following the advice of the 1968 Jerry Lewis film, "Don't Raise the Bridge, Lower the River," MPEG-5 compression presumably offers a way to lessen the flood of bits created by shifting from 2K HD to 4K UHD resolution. Moreover, looking at the 10-year pattern followed by releases of new video compression standards, it might reasonably be expected that yet more

[10]Since the MPEG-3 name was assigned to the aborted attempt to develop a separate standard for HD, the naming of released MPEG standards skips from MPEG-2 to MPEG-4. The MPEG-4 standard was released in 2003, and received a set of major fidelity range extensions (frext) in 2005. These extensions included a new "High" profile that significantly improved compression rates over the original "Main" profile (e.g., by replacing CAVLC entropy coding with CABAC). Therefore, as HD programming started to spread widely after 2006, most broadcasters adopted High profile as the preferred version of MPEG-4.

advanced MPEG-6 compression will arrive around 2023, presumably just in time to handle the next 10X increase in bits, created by shifting from 4K to 8K UHD resolution.

Regrettably, for this optimistic scenario, new compression standards are forced to follow a path of diminishing gains. The logic here is not hard to see. Unlike Moore's Law (mentioned earlier as the prototype for seemingly endless improvements in capacity), compression counts down rather than up. Counting up—in the case of Moore's Law, doubling transistor counts every couple of years—pushes against an elastic ceiling with no apparent limit to its expansion capabilities.

Alas, the logic of counting down is very different. Rather than an indefinitely elastic ceiling above, there is a hard floor below. Recall the two-fold strategy that underlies digital compression. First, remove everything redundant ("lossless" compression). Second, when everything redundant is gone, eliminate everything unimportant ("lossy" compression). When compression has eradicated all the insignificant data (however defined), the limit of what it can do (without comprising quality) has been reached.

Moreover, the closer the approach to the fixed floor set by the pool of all and only significant data, the higher the resistance to further gains. Using relatively modest compute power, MPEG-2 takes 1500 Mbps of raw HD video down to 15 Mbps—a reduction of 99%. By throwing a lot more computational resources at the task of compression, MPEG-4 eliminates a further 67% of the remaining data, leaving a residue of just 5 Mbps.

By throwing an inordinate amount of computational energy at the task, MPEG-5 *might* be able to eliminate another 50% of what remains, reducing 5 Mbps to 2.5 Mbps. (At least, that is the goal set for MPEG-5, though current demonstrations of the technology often settle for reductions of 20-30% rather than the postulated 50%.) MPEG-6, should it be developed, would have to work far harder still to eke out another 25-33%, going below 2 Mbps. And it is likely that no amount whatsoever of computational power could ever push HD video below 1.5 Mbps (setting the hard floor at one significant bit in every thousand).

Note the diminishing returns from additional compression. The vast majority of what can be achieved by compression is done by MPEG-2, eliminating 99 out of every 100 bits. MPEG-4 manages a substantial further gain by tossing out two of every three remaining bits. Realistically speaking, at this point, the bit stream has mostly been wrung dry. Only 1 bit in 300 remains, and it is impossible to eliminate all of this small remainder. Worse, in absolute terms, even if the five remaining Mbps went to zero, the savings would still be only half the ten Mbps reduction realized in the previous step. The figure shown below illustrates this curve of sharply diminishing results for dramatically escalating efforts.

To be sure, the above numbers are merely illustrative and lack both scientific authority and universal validity. But the basic point they make is inescapable. There is hard floor to what data compression can achieve. The closer one gets to that bottom, the less remains to done, and the greater the difficulty of doing it. The future of video compression, therefore, is a set of steadily diminishing gains for rapidly increasing effort. In truth, the last dramatic leap possible with compression technology was MPEG-4. MPEG-5 will work much harder to achieve less, and any future video compression standard will have to work harder still to achieve even less.

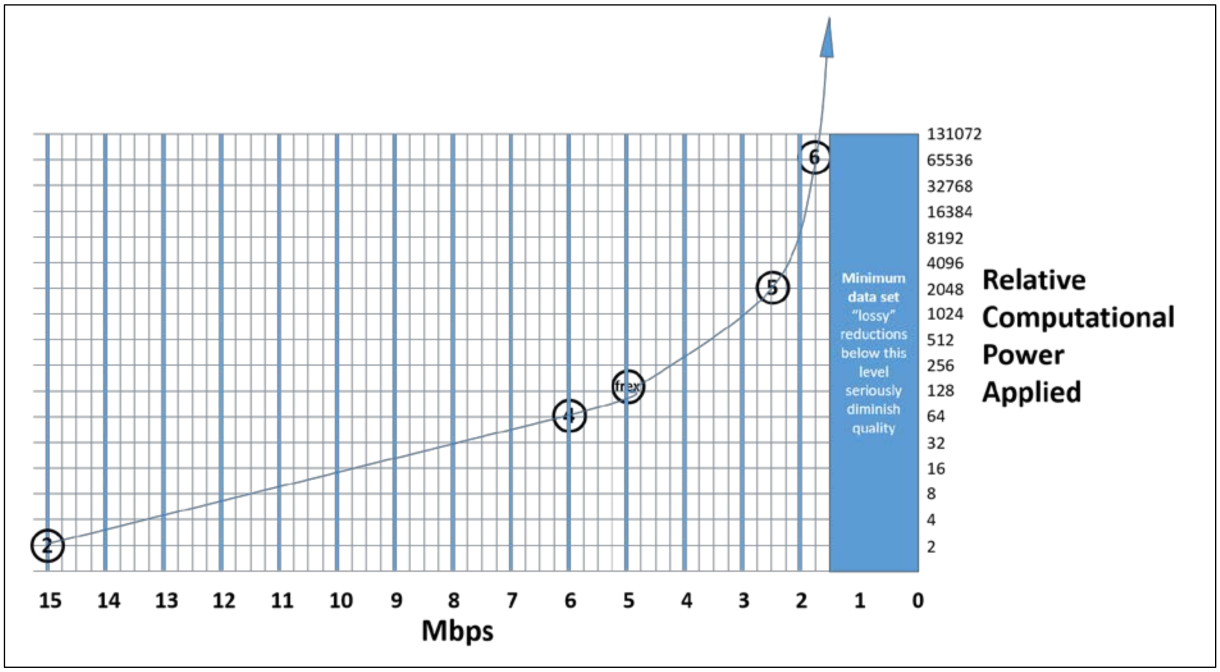

The compression curve for HD video, from MPEG-2 to a hypothetical MPEG-6. As compression approaches the floor set by the amount of non-redundant, non-trivial data (generously, if arbitrarily, fixed here at 1 bit in 1000), real gains decrease while the amount of effort required to make further progress rises exponentially. The assumption used in this graph is that successive MPEG standards appear at 10 year intervals, allowing time for 5 "Moore's Law" doublings in available computational power. This means each new standard has roughly 32X the power of the previous standard to throw at the compression problem, starting from an arbitrary level of 2 for MPEG-2.

5.3.3 4K UHD Compressed Bitrates

Given this logic of diminishing returns, what is the bottom line for 4K video? As we have seen, the increase in bits required by this new UHD resolution standard is a compounding problem, composed of three independent terms: more frames per second, more pixels per frame, and more bits per pixel. When fully realized, this equation multiplies out to an order of magnitude increase for 4K UHD over 2K HD, raising raw data rates from 1.5 to 15 Gbps.

Now apply the above rules for successive compression levels to a raw 15 Gbps UHD data stream. Eliminating 99 of every 100 bits with MPEG-2 technology leaves a staggering 150 Mbps. Eradicating two of every three remaining bits with MPEG-4 reduces this to a still hefty 50 Mbps. MPEG-5—depending on how close it gets to the target of removing one of every two bits left by MPEG-4—will be able to further thin the UHD stream down to 25-35 Mbps.

5.3.4 Modifying UHD Requirements to Lower Bitrates

Looking at the limited headroom left to new compression technologies after MPEG-4, it does not appear that better compression will be of great help in lowering the 10X flood of bits released by 4K UHD resolution, let alone the 100X flood precipitated by 8K UHD. Over time, MPEG-5 may reach its goal of eliminating 1 out of every 2 bits left by MPEG-4 compression. However, it seems a safe bet that no amount of additional compression will ever cut the MPEG-5 bitrate in half again.

Fortunately, there is a simpler way to reduce bandwidth requirements – just cut back on the number of bits generated by new resolution standards. The only truly fixed part of the 4K UHD standard is its quadrupled pixel count. By definition, 4K doubles the W x H dimensions of 2K HD to an array of 3840 x 2160 pixels. So much is inescapable.

There is also substantial pressure to adopt the new HDR10 Media Profile, requiring an upgrade from 24 to 30-bit pixels – a further 25% increase in the data needed. Multiplying out both changes yields a total increase of 4 x 1.25 = 5X the number of HD-bits. The other half of the 10X increase for "full" 4K UHD comes from the requirement to run at double the 30-fps rate of HD. But, since movies have survived for over a century at 24 fps (more or less) without (much) complaint, this part of the new 4K standard probably can be ignored, assuming implementation of both the higher pixel count and heightened color range.

Let us consider, then, a reduced 4K UHD standard, beamed or streamed at 30 fps, requiring a mere 5X increase over 2K HD bit rates. This version of 4K raises the HD raw data rate of 1.5 Gbps to (5 x 1.5) = 7.5 Gbps. Running this 7.5 Gbps 4K figure through the compression ratios set out in the previous part of this series results in the following reductions. First, to 75 Mbps with MPEG-2 (100:1), then to 25 Mbps with MPEG-4 (a further 3:1), and finally to between 12.5 and 16.7 Mbps with MPEG-5 (another 2:1 to 3:2).

Looking at these figures, let's use 18 Mbps as a safe target data rate for 4K UHD. It might be possible to do 4K in less, but a reasonably high-quality version (excluding the doubled frame rate) should be doable within this envelope. As for 8K UHD, let's reduce its required bit rate, to the maximum extent possible, by implementing only its 4X increase in pixel count. This "lite" version of the 8K standard yields a quality MPEG-5 data rate of 4 x 18 = 72 Mbps. And, by applying a theoretical MPEG-6 level of compression, able to use 2023 levels of computing power to generate (say) a further compression of 4:3, this figure might come down to the neighborhood of 55 Mbps.

Thus, by shaving off parts of the new UHD standards, then applying maximum compression pressure to what remains, we have lowered the river of UHD bits as far as seems feasible. The rest of the gain needed to support UHD will have to come from raising the bridge.

5.4 Prospects for More Bandwidth

What are chances of getting, first, 18 Mbps, then 55 Mbps, of bandwidth to a UHD TV? And by when? We will consider this question in two parts: wirelessly and wired. For wireless transmission, we will look at over-the-air broadcasting. For wired transmission, we will look at Internet delivery via cable.

5.4.1 Over-Air Broadcasting

As mentioned earlier, the upgrade from SD to HD resolutions was made possible by shifting TV broadcasts from analog to digital technology. The original digital standard, ATSC 1.0 (from the Advanced Television Systems Committee), was finalized in 1996, the same year the HD standard itself was approved. For purposes of the present discussion, there are two critical points about ATSC 1.0: an 8-VSB RF modulation scheme, and MPEG-2 data compression.

5.4.1.1 ATSC 1.0

8-VSB is an 8-level Vestigial SideBand scheme that, given a 6 MHz RF channel, can transmit 1 of 8 possible 3-bit digital codes (000 to 111) 10.76 million times a second. A little multiplication (3 x 10.76) shows this generates a raw data rate of about 32 Mbps. However, due to overhead[11], the amount of usable data from this scheme is only about 19.4 Mbps.

The selection of MPEG-2, as the technology for video data compression, is certainly unsurprising, since it was the only MPEG compression standard available for broadcast TV in 1996. Which is to say, every built-in ATSC 1.0 RF tuner produced over the past 20 years, from 1996 to 2016 (now), has included a companion MPEG-2 decoder.

Using current ATSC 1.0 technology, then, over-air stations can broadcast HD programs in MPEG-2 at up to 19.4 Mbps. While this exceeds the minimum 15 Mbps bitrate recommended for MPEG-2 HD by CableLabs, it provides little headroom for any higher resolution standard. In particular, a MPEG-2 4K UHD signal would require something like 75 Mbps (far more than even the 32 Mbps raw data rate available).

Over-air 4K programming, therefore, must depend on a new version of ATSC technology. Preferably, the new standard will upgrade both of the critical bandwidth features listed above.

[11]The bulk of the hefty 40% transmission overhead penalty in ATSC 1.0 consists of FEC (Forward Error Correction) codes, which allow receivers to check the bits they receive, determine if any were somehow flipped in transmission, and, if so, correct them.

First, to maximize the number of bits pumped over a 6 MHz RF channel,[12] it will provide a new RF modulation scheme. Second, to minimize the number of bits that need pumping, it will adopt a new compression standard.

5.4.1.2 ATSC 3.0

The good news is that this new standard, ATSC 3.0[13], is already far along in committee. A final version should be approved sometime in 2017. To be sure, judging by the history of ATSC 1.0—adopted in 1996, but not fully implemented until 2011—the prospect of ATSC 3.0 coming to a TV set near you is still some distance off. Setting that issue aside for the moment, however, how much bandwidth improvement does ATSC 3.0 provide?

First, ATSC 3.0 replaces 8-VSB with a more aggressive scheme for bit multiplexing: OFDM (Orthogonal Frequency-Division Multiplexing). This change almost triples the nominal bit rate ceiling, from 32 Mbps to 90 Mbps. Of course, overhead will eat up some of the nominal 90 Mbps bitrate. But even a 40% tax on 90 Mbps would still leave broadcasters with a real data rate of around 55 Mbps.

Second, ATSC 3.0 replaces the 1994 MPEG-2 compression standard with the latest 2013 MPEG-5 compression standard. The good news is that this change makes not only a 30 fps 4K UHD stream possible within the old 8-VSB transmission envelope (at 18 Mbps) but enables 60 fps 4K UHD within the new OFDM envelope (at 36 Mbps).

The bad news is that, unless real OFDM bitrates can be gotten up into the 70 Mbps range (not at all clear), ATSC 3.0 is unlikely to handle even a minimal version of a future 8K UHD upgrade. The good news is that 8K UHD remains a problem for the future rather than a near-term concern.

Regrettably, one could say much the same of ATSC 3.0 itself. Not only is the standard not yet final as of the end of 2016, but serious practical issues render any timeline for its future deployment highly uncertain.

---

[12] Of course, the one thing ATSC 3.0 can't do is increase the 6 MHz RF channel allocated to broadcasters. For better or worse, the radio frequency spectrum is a limited natural resource—a fixed pie that, with the rapid growth of wireless devices of all descriptions, is now beset on all sides by demands for larger slices.

Indeed, as mentioned earlier, with the "2GHz relocation", broadcasters actually surrendered RF spectrum to feed Sprint/Nextel's growing demand for PCS bandwidth. In the face of increasing competition for limited MHz, doubtless the best-case scenario for broadcasters is just hanging onto the 6 MHz over-air channel they now have. As a matter of practical reality, keeping all of this channel for their own use may require over-air broadcasters to upgrade from badly outdated 1996 ATSC 1.0 technology sooner rather than later, simply to show that no one else can put their slice of RF spectrum to better use.

[13] ATSC 2.0, like MPEG-3, was a standard begun but soon overtaken by events, and eventually abandoned without issue.

Chief among these practical obstacles is the incompatibility between ATSC 3.0 and ATSC 1.0. On the transmission side, this means over-air broadcasters will need to make a substantial investment in new modulation, compression, and (to some extent) transmission, editing, control and other equipment before they can deploy ATSC 3.0. On the receiving side, this means none of the roughly 300M TV sets sold in the US since the digital transition began (or many yet to be sold) will be able to receive an ATSC 3.0 signal. In other words, at whatever future cut-over date is selected for the transition to ATSC 3.0, every set that worked up to that moment with 1.0 will immediately go dark. (Unless, by that time, all TVs have dual mode capability, able to run with either an 8-VSB MPEG-2 or OFDM MPEG-5 signal.)

The bottom line here is that, while ATSC 3.0 provides a clear technical solution for broadcasting 4K UHD over-air, it is less evident that it will handle even the minimal demands of a potential future 8K UHD upgrade. Moreover, the incompatibility of ATSC 3.0 with ATSC 1.0 leaves not just its timeline, but even the practicality of deploying it, shrouded in mystery.

5.4.2 Internet Streaming

Unlike over-air bandwidth, which changes in a stepped way as broadcasters switch from one standard to the next, Internet capacity changes in a gradual and fairly regular way as broadband providers make ongoing investments in new and upgraded infrastructure, in response to a steadily rising demand for broadband services. As a result, the Internet has more capacity this year than it did last year, and will have still more next year—and every following year, for some indeterminate (but presumably lengthy) period.

The periodic rise of advertised broadband speeds promoted by the various carriers reflects this increase in average capacity. Currently, for example, Comcast's least expensive plan promises to deliver an "Internet Download Speed" of "up to 25Mbps". Or, for a higher monthly charge, customers can purchase a download speed of "up to 50Mbps". Hence, it appears OTT delivery is already past the 18Mbps needed for minimal MPEG-5 4K UHD video, even in Comcast's lowest service tier, and well past that point in their higher service levels. So, does 4K/8K UHD video pose any problems for OTT Internet broadband delivery? And, if so, what are they?

5.4.2.1 Internet Speeds and Their Rate of Change

The fact that Internet providers advertise broadband speeds as "up to" some number (for example, 25 or 50Mbps) is a bad start on knowing average Internet speeds. The "up to" number listed by a carrier is likely to be a best-case scenario. But, in a world where the only sure thing often seems to be Murphy's Law[14], this "marketing" number is unlikely ever to be realized.

[14]Succinctly stated, Murphy's Law is the principle that, "If anything can go wrong, it will". Which is to say, in the real world, it is generally advisable to plan for the worst-case scenario rather than the best-case scenario.

Formerly, the rule of thumb was simply to divide the listed number by 2, in hopes of getting something like half the advertised rate.

However, the most recent annual report from the FCC ("Measuring Broadband America," December 2015) indicates broadband speeds are up significantly in the past few years. According to the FCC, broadband speeds nearly doubled from 2014 to 2015, with 90% of the customers of the top carriers receiving at least 95% of their advertised rate.

Akamai's "State of the Internet" report for Q2 2016 is also upbeat, if more restrained. Over the past year, it claims peak connections speeds in the US have increased by about one-third, while average connection speeds are up better than one-quarter (28%).

The actual numbers reported by Akamai, though, are relatively modest: a current average peak speed of nearly 70Mbps, but an average actual speed of just over 15Mbps. Since peak rates, by definition, are both rare and fleeting, it seems best to use the average number. And this metric indicates that whatever "up to" tiers broadband providers currently promote, in practice, for the majority of users, at the current rate of improvement, an average speed of even 25Mbps is still at least two years distant.

But, even supposing an average actual broadband speed gain of about 25% a year for the foreseeable future (rounding down Akamai's reported 28% average gain for the last year), the result is still encouraging for broadband 4K UHD delivery. While the current average speed of 15Mbps does not quite reach the 18Mbps threshold set above for minimal 4K UHD delivery, in another year, with another 25% gain, it should pass that mark. And, by 2022, for most broadband users, 18Mbps should appear as nothing more than a steadily receding point in time's rear-view mirror.

Moreover, since the current 15Mbps average must be composed of numbers that fall as much above that mark as below it, some broadband subscribers must be enjoying 4K UHD speeds right now. So, a minority audience is already in place for 4K UHD over broadband. This group will only grow larger year-by-year until, by 2022, it constitutes a substantial majority of viewers.

To be sure, there are many qualifications to this favorable conclusion about the practicality of delivering 4K UHD over the Internet, both now and in the future.

5.4.2.2 Practical Internet Data Delivery

Projecting forward the 2016 Akami number of 15Mbps, growing at a rate of about 25% a year, brings average speeds up to 18Mbps in another year (the minimum threshold we set for quality 4K UHD), and promises to triple that rate by 2022. If we stopped the discussion of broadband video with these numbers, the prospect of streaming 4K UHD over the Internet, if not immediately, at least in the next few years, could hardly look brighter.

But, of course, any claim about Internet speeds is a perfect illustration of the maxim that there are lies, damn lies, and statistics. An average connection speed of 18Mbps in 2017, or even an average speed of 45Mbps in 2022, may turn out to be little or no bandwidth at all in any given case.

The most obvious qualification about Internet connection speeds is that these rates are not tied to either a person or a device, but are rather the connection bandwidth for a *subscriber*, which is typically a whole household of individuals and their devices. As of the most recent Census (2010), the average U.S. household is 2.58 people and (according to Nielsen) contains about 3 TV sets. Let's assume these figures remain roughly accurate today.

Then, if all 2.58 people in an average household gather around just one-third of their TVs, and every other broadband appliance in the house is off, that TV may, in fact, enjoy all 18Mbps that will be that subscriber's average share of the Internet in 2017. But, if they split up into three parties, and each 0.86 person takes one of the three available TVs for his or her use, then that 18Mbps subscriber bandwidth divides likewise into three streams that average 6Mbps each—leaving no one with anything approaching 4K UHD bandwidth. And, even this one-third share of total subscriber bandwidth assumes no other data-hungry appliances—smartphones, tablets, computers, game machines, etc.—are active at the same time, each demanding its share of bandwidth.

Nor does Internet sharing stop at the walls of a household. You may be alone at home with only one TV active. But suppose, when you attempt to stream a 4K UHD version of "House of Cards," many of your fellow subscribers also start streaming 4K. While, at the same time, other subscribers on your local feed start downloading TB files. Then, the likely result is that, at that moment in time, no one on your shared feed will be able to get their fair share of "average" bandwidth.

Moreover, even if there are no local bandwidth issues, either in your household or on your shared feed, there is still the risk of catastrophic failure elsewhere on the Internet. For example, back at the ranch, your OTT video server may be overwhelmed by a sudden surge in demand or (what is much the same thing) a "denial of service" attack. Or your video stream—somewhere between issuing intact from a server and arriving, still intact, at your TV—might fall victim to any of the numerous other ills that can afflict Internet traffic (e.g., failing switches, lost packets, unresponsive name servers).

All of these public network problems—shared usage, bottlenecks, overloads, failing servers, bad switches, cyberattacks, etc.—concern what we might term the physical reality of a global data network. But subscribers also must contend with another class of issues, namely, data throttling or data caps imposed by the service providers themselves. Yes, the same providers that promise speeds of "up to" so many Mbps in their ads.

5.4.2.3 Economic Limits on Data Streaming

The issue of data caps and data throttling by service providers are not so much real physical problems with Internet data delivery as economic constraints placed on these data services by vendors, to manage their costs and provide all subscribers with reasonably uniform levels of access. For example, some carriers are currently rolling out plans that carry a 1 TB monthly data cap, arguing this is an ample limit.

Of course, like the claim that a basketball player is "short," the term "ample" is relative to context. Compared to a cellular plan with a 10 GB monthly limit, a limit of 1 TB a month is more than ample. Compared to streaming an uncompressed 6K UHD file, it is 10 minutes.

5.4.2.4 Streaming 4K UHD Over the Internet

For 4K UHD video, 1 TB/month is somewhere in the middle. A streaming rate of 18 Mbps is 18 MB every 8 seconds, 135 MB in a minute, 8.1 GB in an hour. The 1 TB cap, then, allows about 123.5 hours of minimal 4K UHD video/month. According to Nielsen, the average person watches 5 hours of TV a day. Suppose now that 1 TB plan data is not being frittered away on any other uses (emailing, photo sharing, web browsing, downloading music, etc.). Then, a 1TB cap allows one person in a household to watch 4K UHD video pretty much all they want—for the first 25 days of each month.

5.5 Over-Air vs. Over-Top Delivery

Comparing Over-Air Broadcasting (OAB) with Over-The-Top (OTT) Internet broadband delivery reveals contrasting scenarios of good and bad news.

5.5.1 Advantages and Disadvantages of OAB

For OAB, the good news is the availability of a dedicated 6 MHz channel. The bad news is that's all there is and (in the face of escalating demand for RF bandwidth) all there is ever likely to be. Of course, as ATSC 3.0 proves, improvements to technology over time make it possible to do a lot more with the same fixed resources. However, the rest of the bad news for the long-term future of OAB is that all technological improvements are subject to a law of diminishing returns.

For video compression (as discussed earlier in this series), after discarding everything redundant, irrelevant, and insignificant, a hard floor is reached where all and only significant data remains. Past this point, every additional bit removed by continued compression must come at the price of degraded quality. Similar constraints apply to data multiplexing over a fixed-width channel. Pipe width and light speed ultimately constrain the amount of data a channel can move in a given unit of time. Somewhere before reaching that limit, the cost of cramming additional bits down the pipe must outweigh any benefit conferred by the incremental gain in available data.

ATSC 3.0 is not yet at the ultimate limit of possible improvements to compression or multiplexing, but it is far closer to those limits than was ATSC 1.0. It is certainly close enough to what is possible that no significant further advances are now visible, even on a relatively distant horizon. ATSC 3.0 will have to carry over-air broadcasting to and, very likely, well past 2050. Which is to say, if the ATSC 3.0 envelope can be stretched to accommodate some form of 8K UHD, over-air broadcasting will be able to upgrade from 4K to 8K UHD. Or else, not.

In summary, 1996 ATSC 1.0 technology was sized to deliver HD video and has to be upgraded to ATSC 3.0 to accommodate the much higher bandwidth demands of new 4K UHD. The major issues with this transition are the incompatibility of 3.0 with 1.0 (and the consequent haziness of the timeline for 3.0 deployment) and the still vague upper limits of the 3.0 technology.

These qualifications notwithstanding, it seems likely that, within the next five years or so, ATSC 3.0 will provide a reliable over-air mechanism sized to handle 4K UHD. The future of 8K UHD is less clear, but it is at least possible that it, too, can be managed within the envelope of improving Internet and ATSC technology, at some still more distant time.

But even if ATSC 3.0 justifies all possible technical optimism, there remain serious practical considerations for broadcasters, many of whom are still in the process of fully upgrading from SD to HD. To what extent will 4K UHD require ripping out all their lately acquired HD infrastructure? Enthusiasm for this prospect, likely meager to start, is sure to be depressed still further by the reflection that the reward for successfully negotiating the 2K-to-4K transition may be nothing more than the opportunity to do it all over again, to support a 4K-to-8K transition.

5.5.2 Advantages and Disadvantages of OTT

Turning to OTT broadband delivery, the good news is that, although broadband is subject to the same limits on compression technology that apply everywhere else, there is no hard constraint on broadband pipe size. The size of a wired pipe can always be increased, e.g., by shifting from copper to fiber optic cabling, or by adding more fibers to optical cables.

For broadband, the bad news is that the Internet is a shared rather than a dedicated resource, and (by definition of www) a resource shared world wide. With the Internet, the real constraint on bandwidth is not technology but economics. Practically speaking, a shared world-wide resource cannot be built out to the prohibitively expensive standard of maximum possible demand. Rather, some more affordable standard must be invoked, like “peak” loading.

Peak loading is simply an estimate of the worst-case (highest) actual demand. Statistically speaking, even under the most favorable conditions, this sort of calculation is sure to be wrong some of the time and, under less favorable conditions, may be wrong most of the time. But peak estimates, regardless of their accuracy, are only applicable in constrained circumstances. Thus, Amazon or Netflix may understand the demand for their specific services in a given geography, and design capacity accordingly. But the Internet as a whole is too vast and various for meaningful peak loading estimates.

Instead, we must make do with average available bandwidth, i.e., the bandwidth left after accounting for both real physical issues and imposed economic limits. Some of the time, perhaps even most of the time, the Internet supplies us with all the bandwidth we need (if not all we might desire). By Murphy's Law, however, just at what always seems the most inopportune time, it becomes maddeningly slow.

Indeed, we are protected from Internet apocalypse only by inertia. If any significant shift towards UHD video occurred by everyone at once—not just the Video on Demand (VoD) OTT vendors, but by all the sites that stream video, whether commercial, corporate, public, or personal—the Internet would soon stop working for anyone. The same would happen if a substantial fraction of the audience that receives TV either free over the air (OTA, 17%), or via paid cable/satellite subscription (75%), were to join the still small minority of mostly young Internet TV "cable cutters" (8%).

To repeat, the Internet is a perfect illustration of the claim that there are lies, damn lies, and statistics. Statistics about Internet usage are just that, statistics. Your numbers can and will vary widely, not only from year-to-year as the Internet grows, but from day-to-day, hour-to-hour, and sometimes minute-to-minute.

Still, there is a qualifiedly optimistic moral here. Under some circumstances, for some of the people, some of the time, the Internet either now is or, within a few years, will be capable of delivering a minimal 4K UHD stream. Always assuming, that is, that no sudden surge of demand, significant shift in viewing preferences, or act of global terrorism intervenes to shut the whole thing down for everybody pretty much all of the time.

6 Is 16K SUHD Coming … Someday?

Is that, then, the likely future of resolution technology? Will the coming decades of the 21st century see a continuing series of upgrades to and past 8K UHD, as better technology makes it possible (at a 25% rate of annual improvement) to transmit an order of magnitude more bits every ten years? Frankly, it seems unlikely.

We began this paper with the observation that SD was the only TV resolution standard in North America from 1941 (the original NTSC standard) until 1996 (the 1080i/720p HDTV standard).[15] However, although 55 years elapsed between the initial SD resolution standard and a second, higher bandwidth, HD resolution standard, the third and fourth higher bandwidth standards, 4K and 8K UHD, were both approved in 2012.

This history shows dramatically shortening periods between higher resolution proposals: 55 years (SD to HD), declining to 16 years (HD to 4K), declining to zero years (4K to 8K). The natural question raised by this progression, assuming 4K gets us through the second and perhaps third decades of the 21st century, and 8K can outlast the 2030s, is whether some new 16K/32K super-ultra-high definition (SUHD) standard is destined to appear, say, before the end of the 2040s. And so on.

In short, is there any natural end to this progression of increasing resolutions? Or, with steadily improving technology, do pixels multiply forever, from millions a second with SD to billions a second with HD/UHD, and ultimately to trillions, quadrillions, and more a second with future SUHD standards?

The clear answer to this question, I think, is "no." Potentially unending upgrades will end not because of any physical limit in the ability to transmit bits per second (assuming such a limit exists), but because, with digital technology, it is never about what is *possible*. It is always possible to add more bits to any existing number of bits. Rather, it is about how many bits are *enough* for a given purpose. With digital technology, you simply have to quit when continuing is pointless. Enough is enough.[16]

---

[15]The PAL standard, which emerged in Europe in the 1960s and was widely adopted outside North America, did not change the fundamental SD bandwidth equation established by NTSC. While PAL has 20% higher resolution than NTSC (576 vs. 480 visible scan lines), NTSC has a 20% higher frame rate than PAL (30 vs 25 fps). From the standpoint of data requirements, these two alternatives—moving 20% bigger PAL data sets less often vs. moving smaller NTSC data sets 20% more often—essentially cancel, leaving the amount of data moved per second in PAL and NTSC formats very much the same.

Just as NTSC and PAL are two variations on a single SD resolution standard, 1080i and 720p are best regarded as two variations on a single HD standard, since they, too, are very similar in terms of their overall bandwidth requirements (1080i moves more data less often, while 720p moves less data more often). Note, however, that 4K and 8K UHD, though also grouped together under a single "UHD" label, are two fundamentally different standards in terms of their bandwidth requirements. 8K demands nearly an order-of-magnitude jump in data rates over 4K.

[16] For a deeper discussion of this point, see Appendix 1: Binary Color Coding.

6.1 4K Enough Resolution for a Living Room TV

The real question, then, is how much resolution is enough resolution?[17] For a TV screen, the quickest way to answer this query is to consider display size in light of the powers of human visual resolution. It seems reasonable that the latter capability has to be at the heart of any discussion of how many pixels are enough for purposes of viewing.

To take up the subject of resolution first, the critical fact about human visual resolution is that it drops off in a linear way with viewing distance. Earlier, we mentioned a "retina display" as any screen with a resolution of 300 ppi or higher, since that is the point at which pixels become too small to be individually distinguished by the human eye. The critical qualifier left out of that earlier discussion was viewing distance. For 300 ppi, the assumption is normal reading distance, or about 12 inches—about as close as most people want to get to anything they need to scrutinize.

But TVs, at least in sizes typical for a living room, are not meant for close-up viewing. Say the average TV viewing distance in a US living room is around 9 feet. The fact that visual resolution drops off in a linear way with increasing distance means that, at 9 feet, we can discriminate about 1/9th as many pixels per inch as at 1 foot. Using 300 ppi as the limit of resolution at 1 foot, then 300/9 = 33 ppi is the limit of our normal TV viewing resolution at 9 feet.[18]

Armed with this figure and the knowledge that a full HD screen is an array of 1080 lines of 1920 pixels each, a little math will let us pick out an appropriate HD TV set size, where "appropriate" means the biggest screen with no fewer than 33 ppi. Buying a smaller screen would, of course, increase resolution (ppi), but the tradeoff is bad since none of the resolution gained will be

[17]Consider film restoration as one example of practical digital limits. According to John Lowry, of Lowry Digital (the studio responsible for transferring classic films like North by Northwest, Gone with the Wind, Citizen Kane, Star Wars, Indiana Jones, and James Bond to digital format for DVD distribution), 4K is enough to "capture everything on the film—everything". Whereas higher resolutions are "pointless". Since the amount of digital resolution needed to match traditional analog film resolution is a subject where opinions differ, an extended version of Lowry's quoted statement follows. It should be noted that Lowry is using the terms "2K" and "4K" as these are defined by the DCI (Digital Cinema Initiative) standards—2048 x 1080 and 4096 x 2160, respectively. While the number of lines are the same, the DCI standards have roughly 7% more pixels per line than the HD/UHD standards.

"*If I'm going to restore a film, my objective is to capture everything that is on that negative, which probably has a limit somewhere in the 3- to 4K range…. If you scan it at … 2K, there's all kinds of information on that film that you just haven't got. If you scan at 4K, it captures everything on the film—everything…. We create a digital master … that is just as good as the original camera negative in terms of resolution and grain structure…. In all the measurements I've done, I've yet to see much information on a film right up there at the 4K level—it usually rolls off between 3 and 4K. We've experimented at 6K, but, frankly, it's pointless on a standard 35mm film frame.*" (quoted in "Creating the Video Future" by Josef Krebs, Sound & Vision, November 2004, pp 110-112.)

[18]This rule of thumb (300 divided by feet) scales amazingly well. At 300 feet (100 yards, the length of a football field), the limit of human resolution is a square about 1 inch on a side.

noticeable from a distance of 9 feet, while the loss of viewing area will significantly diminish the overall immersion of the viewing experience.

So, what is the right size when choosing an HDTV for a living room with a typical 9-foot viewing distance? The short answer is a 65” set. A TV screen with a 65” diagonal measurement has a width of about 57” or a pixel density of 1920/57 = about 34ppi, or a shade more than we can discriminate at 9 feet. This screen will show us every detail we could conceivably make out.[19]

Now assume that, when purchasing a new TV, it is not feasible to also change the viewing distance. In this case, the maximally useful resolution figure will remain fixed (e.g., at 9 feet, 33 ppi). Then, the result of upscaling from HD to 4K UHD will be to double screen width. For example, a 4K UHD format screen with the same 34 ppi resolution as a 65” HDTV would measure 130” diagonally, and be about 113” wide by 64” high.

Suppose we can surmount the challenge of placing a screen more than 9 feet wide and 5 feet high in a typical US living room. Then, if the nearly 5-foot wide 65” HDTV did not provide a sufficiently immersive viewing experience, a 130” 4K UHDTV, with a screen approaching 10 feet wide, might reasonably be expected to do so. And, even if it does not, doubling screen size again, with an 8K UHDTV, does not seem feasible in any but the most palatial homes—since it requires a living room with 11-foot ceilings and a wall uninterrupted by doors or windows for 19 feet.

Thus, if the question is: “How many pixels are enough for a large screen TV designed for the average US living room?”, the answer is 4K. 8K appears to move us from the realm where “enough is enough” into the realm of serious overkill.

---

[19]A 70” screen, with a width of about 61”, would render a bit less detail (1920/61 = about 31 ppi) than the 33 ppi we could, in principle, discriminate at a distance of 9 feet. Though, in reality, there is no point in being too precise about any of the numbers regarding human vision, ppi values, screen sizes, or seating distances.

Human vision, in fact, varies widely between persons (all the way from extreme near-sightedness to extreme far sightedness). It also changes considerably over time for individuals, and is affected by a variety of widely variable environmental factors, including lighting and atmospheric conditions like haze or smoke. As a result, the ability to discriminate two dots at a given distance may be very different not only for different people, but for the same person at different times, and for everyone at the same time as environmental factors fluctuate.

As for seating distances, simple changes in position (e.g., sitting up vs. leaning forward or back) can easily move the associated viewing distance by a foot or more.

The bottom line is that, except where otherwise indicated, all the base numbers in this paper should be regarded as illustrative or typical, since absolute precision about human capabilities (etc.) is generally unobtainable. Likewise, unless otherwise indicated, computed figures should be regarded as approximate. Many have been worked out on the basis of starting values, often arbitrarily plucked for purposes of illustration from a range of possibilities, then manipulated by rounded and simplified calculations. For this reason, they are unlikely to be exactly the same as equivalent numbers found elsewhere. Hopefully, they remain similar enough to other, independently derived figures, to be regarded as correct within the wide limits of error imposed by the variable nature of the subject matter.

6.2 What About 8K (and Higher) Resolutions?

8K resolution allows an "optimal" viewing distance of 9 feet in front of a screen 11 feet high and 19 feet wide. These sorts of H x W dimension surpass the merely "immersive," on the way to a truly overwhelming experience for home consumption. Besides which, they seem impractical in all but exceptional circumstances, based simply on the dimensions of ceiling height and wall width available in most homes—independent of any new display technologies, e.g., "wallpaper" OLED, that would allow future home TVs of any size to be unrolled directly onto a wall.

Thus, if 8K UHD falls beyond the limit of what seems feasible for home viewing, it seems safe to assume the madness of constantly escalating resolution standards must end with 8K. We should not have to worry about 16K SUHD (15360 x 8640) TVs, with their demands for 22-foot ceilings and 40-foot walls to support a "typical" 9-foot viewing distance. Even shrinking a 16K screen by moving closer to it—reducing pixel size/spacing all the way down to a 300 ppi "retina" viewing distance of 1 foot—does not help (enough). It still leaves a "tablet" screen of enormous size, over 4 feet wide and nearly 2 and a half feet high—far too big for any practical handheld use.

Thus, I conclude, we are safe from 16K or greater screen resolutions and, in all but niche markets, safe from any real need for 8K UHD as well. Which is not to say it is impossible for screen manufacturers and others to attempt to generate a widely-perceived need for 8K UHD. Indeed, we fronted this paper with an earnest plea to upscale from 4K to 8K UHD as rapidly as possible.

Let us assume, then, as the market for 4K TV sets becomes saturated, the future holds a rising tide of opinion promoting the alleged advantages of 8K resolution (with the quote displayed at the head of this paper merely a harbinger). When faced with this sort of promotion, just keep in mind that (ignoring any changes in color depth or frame rate), the only difference between HD and 4K, and between 4K and 8K, is a 4X increase in the number of pixels displayed.

If screen size is held constant, the only way to accommodate this change is to cut in half the distance both between rows of pixels and pixels in a row. This new spacing, in turn, halves the "recommended viewing distance" for the TV set. Conversely, if pixel spacing and its associated "recommended viewing distance" are held constant, the only way to accommodate this change is to make TV screens twice as wide and twice as high. Either way, the screen will loom four times larger in the visual field, generating a more immersive experience.

The simplest way to summarize this increasing “loom” factor of higher resolutions is by showing the decreasing series of fixed ratios between screen width[20] and viewing distance for increasing resolution standards. For HD, the fixed ratio for “optimal” viewing distance is 1.8X the screen width. Thus, for example, take the earlier computation of a 65” HD screen as the perfect size for a room with a viewing distance of nine feet. Nine feet is just about 1.8X the roughly 57” width of a 65” HDTV screen.

Obviously, making the screen half the size and viewing it from half the distance (or any other proportional adjustments in these dimensions), does not affect this ratio. So, the ratio of 1.8X screen width as the ideal viewing distance for HDTV holds constant for all HDTV screen sizes.

Since 4K supports the same resolution with a screen twice as wide, the shift from HD to 4K halves the HD ratio number, i.e., the ideal viewing ratio for a 4K UHD screen is 0.9X the width of its screen. Upgrading from 4K to 8K would cut this ratio number in half again, moving the ideal viewing distance for 8K UHD to 0.45X its screen width.

The figure below illustrates this set of decreasing ratios. The shift of viewpoint from far outside to far inside the “TV Square” (shown in blue) graphically demonstrates the increasing “loom” factor of wider TVs that results from progressing through HD to 4K and on to 8K.

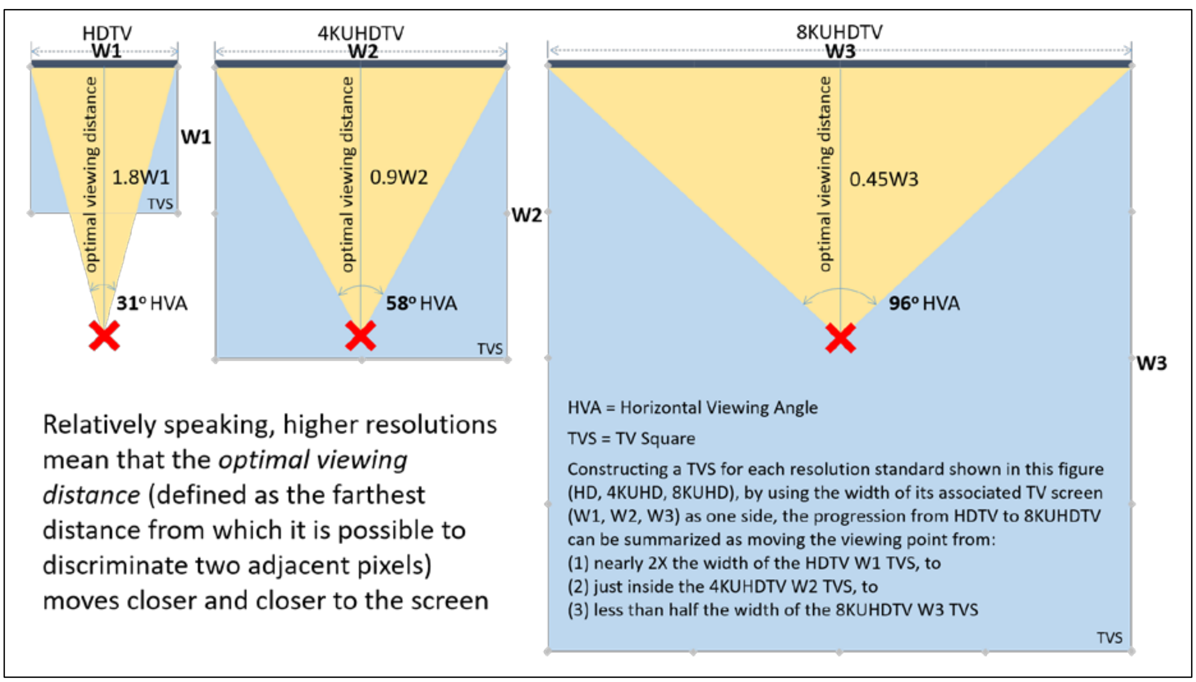

[20]Note the relevant screen dimension here is width, not the commonly quoted diagonal measurement. For 16:9 HD/UHD screen ratios, width can be estimated by subtracting 13% from the diagonal dimension. 13% is a bit awkward for direct mental calculation, though it can be factored into 10% plus a short third of 10%. For example, the width of a 65” screen, measured diagonally, can be guesstimated by subtracting 6.5 inches (10% of 65), plus another 2 inches (the down round of 1/3 of 6.5) = 8.5 inches, from 65. This answer, 65 – 8.5 = 56.5 inches is close enough to the true value of 56.65241 inches as to make no practical difference.

For reference, commercial cinemas built to THX specifications have a minimum horizontal viewing angle of 36 degrees (from the last row) and a maximum viewing angle of 62 degrees (from the first row). The "sweet spot" for theater viewing falls in the middle rows, at a horizontal viewing angle somewhere around 45-50 degrees.

Note the jump from HDTV to 4KUHDTV takes us all the way from 31^o, just outside the lowest recommended 36^o "back row" angle for a commercial theater, to 58^o, just under the highest recommended 62^o "front row" angle for a commercial theater. Not to mention, significantly closer to the screen than the 45-50^o viewing "sweet spot." Again, the moral seems to be that 8KUHDTV, with its overwhelming 96^o horizontal viewing angle, is a bridge too far for home use.[21]

7 Summary

Since this is a long paper that covers a broad range of topics, let's conclude by summarizing the major points made.

1. The primary purpose of a digital resolution standard is to fix pixel count, specified as a W x H array of rows and columns. A digital resolution standard also specifies the number of bits that compose the data for a pixel, and the video frame rate.
    1. The "full HD" standard is a pixel array 1920 wide by 1080 high, based on "8-bit" color with 24-bit (3 x 8) RGB pixels, shown at a rate of 30fps. Multiplied out (1920 x 1080 x 24 x 30), this is a data rate of about 1.5 billion bits per second.
    2. The "4K" UHD standard doubles both width and height dimensions of the full HD pixel array, to 3840 x 2160, based on "10-bit" color with 30-bit (3 x 10) RGB pixels, shown at a rate of 60fps. Multiplied out (3840 x 2160 x 30 x 60), this is a data rate of about 15 billion bits per second.

---

[21]The extravagance of 8KUHDTV can be approached from the other end, by looking at the limits of human vision. While no horizontal viewing angle of a flat surface can reach 180 degrees (at least, as long as there is any separation of viewpoint from surface viewed), the human visual field is actually slightly wider than a half-circle, at about 190 horizontal degrees. The binocular field, however, defined as the range of points that can be seen by both eyes at the same time, is substantially narrower, about 120 degrees.

The remaining 35 degrees on each side of our central binocular field constitutes our peripheral fields of vision. By definition, since objects in our peripheral fields can be viewed by only one eye, peripheral vision lacks stereoscopic perception and its associated sense of depth.

At 96 degrees, 8KUHDTV underlaps our entire range of binocular vision by only a 12-degree sliver on each side. A potential future 16KSUHDTV, with an optimal viewing ratio of just 0.225 screen width (half the 0.45 ratio of 8K), would provide a truly extreme horizontal viewing angle of about 132 degrees, actually exceeding our binocular field and overlapping into our peripheral vision.

   3. The "8K" UHD standard doubles both width and height dimensions of the 4K UHD pixel array, to 7680 x 4320, based on "12-bit" color with 36-bit (3 x 12) RGB pixels, shown at a rate of 120fps. Multiplied out (7680 x 4320 x 36 x 120), this is a data rate that approaches 150 billion bits per second.
2. Everything about pixels beyond their W x H array, their length (color coding scheme), and the rate at which they succeed each other falls outside a resolution specification. Specifically, pixel size, screen metrics like pixels-per-inch (ppi), and even how bright and vivid an RGB color code appears on a screen, is a function of the display used to render pixels.
3. Any assessment of the new UHD 4K (and 8K) resolution standard must be done in two parts.
   1. One part assesses the impact of the new standard on consumers, i.e., the value of a new resolution standard. More specifically, what enhancement does UHD bring to the viewing experience, and what are the challenges consumers face when acquiring new UHD-capable sets? This demand-side assessment is fundamental since commercial products exist to meet demand. Little or no demand means little or no product; high demand results in vigorous efforts to supply product.
   2. The other part assesses the impact of the new standard on businesses involved in the creation, acquisition, and distribution of television programs. Assuming a demand exists for 4K programming (fueled by the widespread availability of affordable 4K screens and their consequent spread among consumers), what are the practical issues confronting businesses involved in the creation, acquisition, and distribution of video, for making 4K UHD programs available to a mass TV audience?
4. The value of HD over SD was sufficiently evident that there was little debate over the desirability of this transition for the TV audience. The long 30-year wait for HD to come into general use, from the first serious proposal about upgrading SD resolution, was a supply-side problem. The entire TV industry had to execute a fundamental technological shift, from analog to digital video, to make HD signals generally available. Creating this new digital technology and its associated infrastructure took considerable time.
5. Now that the digital transition is over, however, no similar technical shift exists (or is needed) to enable another resolution upgrade, from HD to UHD. The difference between HD and UHD programs is just the difference between some bits and an order of magnitude more bits.
6. How this flood of new UHD bits affects the feasibility of doing business depends on where a company stands in the chain of creating, transmitting, and receiving these bits.
7. To begin at the receiving end of the content chain, for TV set makers—once they overcome the challenge of manufacturing 4K screens—4K is all opportunity. UHD creates a whole new market for display sales, as buyers seek to replace their old HD sets with new UHD sets.
8. At the other end of the content chain, the cost of creating programs rises substantially with 4K (given the cost of replacing all the old HD production equipment). However, these added production costs are minor in the overall budget for a major TV show. If 4K production increases audience share, producers will likely regard its added costs as money well spent.
9. OTT (Over-The-Top) companies, who make video content available for Internet streaming, sit one step below content creators. (Although, this split, between creation and distribution, is not

as sharp as previously since major OTT companies like Netflix, Amazon, and Hulu are now also creating original content.)

10. Whether OTT 4K content is generated internally or acquired from external sources, the strength of the OTT position for 4K distribution is that their responsibility ends with making content available in a variety of streaming formats, ranging from high-resolution, high-bitrate 4K versions (if available), to low-resolution, low-bitrate mobile formats. Which version of any content gets streamed to a customer is strictly a function of (1) the receiving device the customer has, and (2) the amount of Internet bandwidth available between sender and receiver when streaming occurs—two issues over which an OTT company has neither control nor responsibility.
11. The problem posed by 4K for businesses, then, does not lie at either end of the content chain, with video reception or video production/acquisition. The real challenge is in the middle, with the distribution of 4K content to 4K receivers. The barrier to an industry-wide transition from HD to UHD programming is the herculean task of transmitting an order of magnitude more bits across available wired and wireless channels in real time.
12. In general, there are two ways to cope with the overwhelming threat of a sudden 10X rise in bits per second, generated by the replacement of current HD content with 4K UHD content.
    1. Eliminate as many of the additional bits as possible, either by employing better compression technology or by implementing only some parts of the UHD specification (a UHD “lite” format).
    2. When all the disposable bits are gone, the only remaining solution is to increase transmission bandwidths sufficiently to enable delivery of those bits deemed indispensable.
13. Better compression is of limited help. Successively higher compression standards improve by progressively whittling away at the pool of unnecessary and unimportant bits in a bit stream. But the more successful each new standard is at identifying and eliminating the bits in this pool, the less there is to do. MPEG-2 is a massive improvement over no data compression at all. MPEG-4 is a big jump in compression efficiency over MPEG-2. MPEG-5 must work much harder to eke out a smaller gain over MPEG-4. And so on, until, at last, no further reductions from compression are possible, because the compression force applied (MPEG-6? MPEG-7?) either leaves no unimportant bits or these bits are sufficiently few and elusive that the effort of hunting them down is not worthwhile.
14. An alternative (and far easier) way of eliminating the bit flood created by UHD is not to implement the less visible but more expensive parts of the standard, e.g., for 4K, the demand to double the frame rate from 30 to 60 fps. Continuing to transmit 4K at 30 fps would, at a stroke, cut the raw data rate increase for 4K in half.
15. Counting everything—including the proposed ability of MPEG-5 to cut MPEG-4 data rates in half, and ignoring any increase for doubled frame rates—it seems that quality 4K frames could be transmitted at a data rate conservatively set at 18Mbps. Even lower rates may be feasible. But data rates approaching MPEG-4 HD (~5Mbps) are out of the question for 4K.

16. What are the prospects for an increase in transmission bandwidths to 18 Mbps? This issue is considered in two parts, first, for wireless Over-Air Broadcasting (OAB), then for wired Internet streaming or Over-The-Top (OTT) broadband delivery.
17. For OAB, the new ATSC 3.0 standard includes changes to improve both video compression (from MPEG-2 to MPEG-5) and data multiplexing (from 8-VSB to OFDM). These two changes together should provide ample bandwidth, not just for 18Mpbs 30 fps 4K video, but for full 60 fps 4K transmission at 36Mbps or higher. It is possible but less sure that ATSC 3.0 technology will also prove sufficient for 8K transmission (another near 10X jump in raw data rates over 4K), though probably only in some limited form.
18. The deployment of ATSC 3.0, however, remains shrouded in mystery. In part, this is simply because the standard is not yet complete (as of the start of 2017). In part, this is because ATSC 3.0 is incompatible with ATSC 1.0—an awkward fact that makes it unclear exactly how or when to implement a transition from 1.0 to 3.0.
19. For Internet delivery of 4K content, the mere existence of 4K receivers, and 4K content queued for delivery to those receivers by OTT providers is inadequate to trigger 4K delivery. The missing parameter is bandwidth. On today's Internet, the 18Mbps required for 4K video transmission is probably more nearly the exception than the rule. At the current growth rate of 25%/year, however, by 2022, today's average bandwidth of 15Mbps should triple to around 45Mbps. That increase should make it possible to accommodate much, if not all of the streaming that takes place in HD today, in a 4X more voluminous 4K format.
20. I conclude the technical barriers to delivering 4K video streams in real time, while substantial, are not insurmountable, either for OAB or Internet streaming.
21. Whether today's TV industry can adjust to 4K UHD programming, though, is not the final question. The more interesting question concerns the long-term future of resolution standards, specifically: 'Is there any end to this road of increasing TV resolutions?'. We have now progressed from SD to HD, and are working on a further transition to 4K UHD. But there is already a still higher resolution standard, 8K UHD, waiting in the wings. As technical progress on the supply side continues to improve through the 2020s into the 2030s, is 4K UHD inevitably going to be followed by 8K UHD? And what about potential successors to 8K in the even more distant future? With continued progress through the 2040s into the 2050s, will 8K be succeeded by 16K? And so on, until, by 2100, the world is at some presently inconceivable 128K resolution standard with data rates of 1.5 quadrillion bps?
22. For the sake of discussion, let us suppose there are no technical barriers to still higher bit rates on the supply side (or at least none that time and human ingenuity cannot surmount). The good news here—at least, for those who fear the ongoing financial and other strains of an endless upgrade cycle—is that technical capability alone is not sufficient to trigger a resolution upgrade. There also must be some value to the potential upgrade.
    1. Consider microprocessors as an example. They transitioned from 4-bits (1971) to 8-bits (1972) to 16-bits (1979) to 32-bits (1985) to 64-bits (1991), all in relatively short order. However, there has been no further transition to 128-bit processors—despite the lapse of 3 decades since the first 64-bit microprocessor appeared, and the relative ease of implementing yet another bit doubling with today's enormously improved

semiconductor technology. Upgrades ended at 64-bit machines because 64-bits are adequate for even the most advanced supercomputer uses.

23. Thus, the more important question for any possible future upgrade in resolution standards is not ‘Is it possible?’ but rather ‘What benefit does yet another resolution upgrade bring to the TV audience?’. If there is no reason for consumers to want higher-than-4K resolution, there is no reason for suppliers to tackle the daunting problem of enabling still higher resolutions.
24. The benefit of a new resolution standard for consumers—the improved viewing experience it generates—primarily results from its increased pixel count and, secondarily, from improvements to color coding for pixels. For viewers, the benefit of increasing frame rate is minor. Not only have movies survived happily for a century at 24fps, but most TV viewers seem to prefer 1080i HD resolution at 30fps to 720p HD resolution at 60fps.
25. Just as HD was a definite boon over SD for TV viewers—once accustomed to HD images, few are willing to watch SD again on a regular basis—objective analysis indicates that 4K is a significant upgrade over HD for viewers. Both its quadrupled resolution and its use of 10-bit color substantially enhance the viewing experience.
    1. The quadrupled resolution of 4K allows for far more immersion in the content. Current HDTV puts the home viewing experience past the back row of a typical movie theater. 4K, by contrast, enables a near-front row experience.
    2. Similarly, the shift from 8-bit to 10-bit color allows use of High Dynamic Range (HDR) color technology, which much more accurately reflects the high dynamic range of the human eye itself.
26. Given the existence of affordable 4K sets, and the technical feasibility of quadrupling bit rates to support minimal 4K programming, the advantages of 4K over HD for viewers, therefore, should be sufficient to make it the standard for homes during the 2020s.
27. However, the same objective analysis argues that resolution upgrades past 4K—including the on-the-books 8K UHD standard—are unnecessary (like processor upgrades past 64-bits). Either they contribute nothing to the viewing experience, or their contribution is negative.
    1. Given the HD-to-4K upgrade moves TV viewers, in movie theater terms, to a near-front-row perspective, any further increase in pixel count seems superfluous. Indeed, quadrupling 4K resolution with the 8K standard—in movie theater terms, moving the viewing perspective far beyond the front row—appears to bring the immersive experience well past the comfort zone of almost everyone.
        i. A related, purely physical problem exists. While 4K screens can be big (e.g., 9 feet wide by 5 feet high) their size is still manageable. However, useful 8K screens are potentially enormous, far exceeding the available wall space of an ordinary living room.
    2. As a rule, for a viewer, binary codes longer than 8-bits for primary colors are not useful. 8-bit “true” color encodes over 16 million RGB colors, far more than even the keenest human eye can discriminate. So-called “deep” color codes of 10-bits and more do not augment the set of RGB colors we can discriminate. Rather, they multiply codes for what, to the eye, are indistinguishable shades of the same color.

   i. An exception to this rule needs to be made for 10-bit HDR color in the 4K standard since HDR does provide a visible enhancement to displayed colors. But 12-bit color in the 8K standard is just billions of more codes for colors no human eye will ever discern.

28. I conclude the added cost and expense of 8K HDR resolution is of no real interest to the home viewer, meaning that the upgrade path from SD to HD to 4K UHD should end with 4K, rather than continue indefinitely on to 8K, 16K, and beyond.
29. If there is a bottom line regarding resolution standards for digital video, it is simply this. The digital revolution—which, at its core, is about converting everything to bits—is never about whether it is, in fact, possible to have still more bits. However long a binary number may be, making it longer is always possible. For digital video, it is possible to increase pixels of 30 bits to 36 bits, or 42 bits, or more. And arrays of pixels (which are nothing but bundles of color bits) can be expanded from 8 million per frame to 32 million, or 128 million, or more. A bit stream composed of successive frames can be pumped up from 15 billion per second to 150 billion, or 1.5 trillion, or more. The real question for any digital technology, then, is not, 'How many bits can we have?', but rather 'How many bits are enough (for the purpose)?' Just as, for microprocessors, the answer to the latter question is 64 bits, for TV viewing, it appears to be 4K UHD.

Appendix 1:
Binary Color Coding

Although the long strings of 1s and 0s that comprise binary numbers can seem quite daunting when first encountered, understanding binary numbering is easy. The basic rule is just that every bit added to a binary number doubles the number of possible combinations supported. This fact can be seen most readily by starting at the beginning, with 1 bit, which has only two possible values (0, 1). Adding a second bit allows four possible values (00, 01, 10, 11). Three bits have eight possible values (000, 001, 010, 011, 100, 101, 110, 111), Four bits have 16 possible values, five bits 32 possible values, etc. By the time you reach the 8-bit values used in “true” color RGB encoding, this doubling algorithm has passed by 64 (six bits) and 128 (seven bits) to reach 256 possible combinations of 1s and 0s between 00000000 and 11111111.

Simply put, the reason adding a bit doubles the previous set of numbers is that it allows writing all the numbers of the previous set twice over. The first time, add a 0 in front of all the previous numbers. The second time, prefix a 1. Thus, for example, compare the 1-bit vs. 2-bit and 2-bit vs. 3-bit values shown above.[22]

Since there are 256 8-bit binary numbers, an 8-bit channel provides encoding for (i.e., can assign a distinct binary number to) 256 different shades of a color. A 24-bit RGB pixel encodes 256 shades each of Red, Green, and Blue, or 256 x 256 x 256 = 16,777,216 “mixed” colors. Although no one can say exactly where the limits of human perception fall, a ballpark estimate of one limit is that the human eye can distinguish, at most, about 150 shades of a color. Accepting this number, in turn, means that, while 7-bit color channels (128 shades) are not quite good enough to express the full range of human vision, 8-bit channels (256 shades) are more than adequate.

Thus, 8-bit channel encoding is called “true” color as the first and simplest color scheme able to capture and reproduce all the RGB colors human beings can see. The term stands in contrast to earlier, simpler, and less adequate schemes, e.g., the 5/6/5 RGB 16-bit pixels of “high” color (which provides a much reduced if still extensive palette of 32 x 64 x 32 = 65,536 mixed colors).

The easiest way to see the “true” color nature of 8-bit color encoding, is to build a color bar out of 256 equal width stripes, each composed of an adjacent shade of, for example, red. Code the first strip as pure red (255/0/0) and the last strip as pure black (no color, 0/0/0). The result will be that this bar, which descends by 256 discrete levels from red to back, does not appear to the

[22] The same rule applies to any numbering system. The reason adding a place in decimal multiples the set of available numbers by 10 is that the additional place makes it possible to write the previous set 10 times over, each time affixing a different base number, 0 to 9, as a prefix.

eye as 256 distinct stripes, but rather as a single continuous gradient, shading from red to black by insensible steps. This experiment shows that dividing a color into as many as 256 distinct shades moves adjacent stripes below the threshold of just noticeable differences. No one can tell stripe one from stripe two, two from three, and so on down the entire length of the 256 strips that compose the color bar.

In fact, for many people, the same would be true of a color bar built up from 128 strips (7-bit channels), but a keen eye under ideal conditions can see very faint stripes in this bar. So, 7-bit color channels are not quite past the limit of human perception. But 8-bit channels, with their associated 256 shades, are well past it. With 8-bit color channels (24-bit pixels), it is possible to code not only all the colors anyone might ever be able to distinguish under any circumstances but many millions more besides that no one can tell apart.

But if 8-bit "true" color channels are sufficient to generate not only all perceivable RGB colors but many millions more besides, it is only natural to wonder 'What is the advantage of a new 10-bit "deep" color standard?'.

Before answering that question, let's consider the binary logic of 10 bits. That much, at least, is straightforward. Since each new bit added to a binary number doubles the number of combinations previously possible, the move from 8- to 10-bit color first doubles the number of possible primary color shades from 256 to 512 (at 9 bits), and then redoubles that number from 512 to 1024 (at 10 bits). 1024 x 1024 x 1024 multiplies out to over a billion (1,073,741,824) possible RGB "mixed" colors, or 64X (4 x 4 x 4) the number of 8-bit colors.

But, to repeat, what is the advantage of moving from an 8-bit "true" color channel, able to code 256 shades of each primary color, to a 10-bit "deep" color channel that can code 1,024 shades of each? Compared to 24-bit pixels, don't 30-bit pixels just add more than a thousand million mixed colors that no one will ever see?

There are two answers to this question. Regrettably, the first is that the marketing of "deep" color often relies on a simple "more is better" argument. If 8-bit "true color" is good, then 10-bit "deep" color must be even better. The name says it all. With 10-bit coding, colors will have to be somehow "deeper." Why settle for a mere 16 million colors when you could have a billion?

The fallacy of this argument is not hard to spot. In fact, 10-bits does not make any 8-bit color "deeper." Its only function is to divide each of the 256 8-bit shades into four separate sub-shades, all of them indistinguishable to the human eye. This point is worth dwelling on because there is a hidden truth lurking here, just below the surface.

In fact, the term "deep" color is a generic term, embracing any color depth past the 8-bit "true" color standard, including 12, 14, and even 16-bit colors. So, if deeper is better, there is no reason to stop at 10-bit channels. Rather, we should move on to add bits to color channels as rapidly as technically feasible. Why not 64-bit channels? Or 1024-bit channels? Indeed, why should the ambition for more color bits ever end?

The lurking hidden truth is that, with digital coding, “more is better” is ***not*** a good argument. It implies code lengths should be expanded indefinitely, always to the limits of current technical feasibility. This claim is just silly. The real question when it comes to digital code lengths is not: Can we make it bigger? That question is not even meaningful since its answer is always “yes.” The real question that needs asking is: How big does a code have to be? Or, alternatively: How big is big enough for the code’s intended purpose?

An example may help. The first microprocessors were just 4-bit machines. But, while 4-bits is O.K. for simple tasks, over the next 25 years, microprocessors expanded to 8 bits, then 16 bits, then 32 bits, and finally to 64 bits. But, since reaching 64-bits over 20 years ago, there has been no clamor for yet another upgrade to 128 bits, and no such demand seems to be on the horizon. The reality is that 64 bits are not only enough for any ordinary computing job, but they are also sufficient for any supercomputing task. Throwing vast resources at creating a whole new generation of unneeded 128-bit machines would be insanity.

Just as microprocessor word length increased from 4 to 64 bits, there is no problem with expanding color channel length. Instead of millions, we could have billions, trillions or quadrillions of colors. It only takes more bits. But none of this augmentation changes the human eye. If the question is: ‘How many bits are needed to exceed the capacity of the human eye?’ – The answer is eight. Add as many colors as you like past the 8-bit “true” color RBG space, it remains certain that no human will ever discriminate even one of these possibilities.

Therefore, let us agree that “more is better” is a bad argument for expanding past 8-bit color. Let us also agree that any such expansion of digital color codes adds absolutely nothing to the range of human color perception. These arguments just raise another question: Why does the UHD standard propose 10-bit color? Surely the experts behind this standard must have *some* good reason for wanting to increase color channel depth.

In fact, there are two sensible reasons for wanting more than 8-bits/256 shades per primary color. The first reason pertains to capture devices, like cameras and scanners. As mentioned when discussing digital compression, the technologies used for image capture are often far more sensitive to different wavelengths of light than the human eye. The fact that the eye tops out at perhaps 150 or so color shades does not limit our image sensors, which may discriminate far more than that number. For these superhumanly sensitive capture instruments, recording all of the information they make available could require 10 bits, 12 bits, or still deeper color channels. Similarly, old color photos contain subtleties the eye does not see. Hence, to digitize all the information available in a family photo album, it’s a good idea to shop for a color scanner that advertises, for example, “48-bit color” (16-bit channels).

The second reason has to do with the "lossy" nature of digital image compression. Since image compression works by throwing away information, if a picture is to undergo several rounds of editing and compression before reaching final form, it's useful to begin this process with a lot more information than is needed at its end. Remember that 8-bit "true" color stops at the very first bit past human perceptual limits. But if the result of editing the millions of colors in an 8-bit image is to reduce them to mere thousands of colors, the losses very likely will be perceptible. In this case, it's a good idea to begin, not with millions, but rather billions of colors, since reducing billions to mere millions is unlikely to result in any perceptible degradation.

Neither of these arguments, however—which have to do with image capture and image processing, respectively—provides any reason for transmitting 10-bit "deep" color to displays. Displays are all about human perception and that, to repeat, stops at 8 bits. Even so, there is a sound reason for providing 10-bit color to displays. But that reason is ***not*** about showing new mixed RGB colors that no eye will ever see.

Rather, it has to do with a technology known as High Dynamic Range (HDR) color. HDR color is not about adding a billion or so colors we can't possibly appreciate, but about overcoming a limitation of our image capture technologies to generate pictures that more closely resemble what we can see.

We have already mentioned one way in which artificial capture technologies may exceed the capabilities of the human eye, namely, in their ability to detect wavelength differences within the visible spectrum too subtle to be noticed by the roughly 6 million color-sensitive cones in our eyes. However, there is also an important way in which the human eye is superior to artificial capture mechanisms. The 120 million or so highly sensitive rods in our eyes, which detect light levels, respond to a much wider range of luminosities than our light-recording mechanisms can capture. This issue is familiar to anyone who has selected a camera exposure value ("f-stop" and shutter speed) for a scene with strong light-to-dark contrasts. The eye, with its naturally high dynamic range, can make out details at both ends of this spectrum. Relatively speaking, our light-recording mechanisms typically have low dynamic ranges.

In practice, this means that setting a camera at one end of its exposure scale results in the darker parts of the image taking on enough contrast to distinguish features, while the lighter parts wash out. Conversely, setting the exposure at the other end of the scale results in the lighter parts taking on enough contrast to make out features, while darker parts black out. Whereas, the compromise of setting the exposure in the middle, results in loss of detail in both the darkest and lightest parts of the image.

HDR color solves this problem, in effect, by taking three different exposures of every scene, one optimized for the lighter parts, one for the darker parts, and one for the middle parts. The best parts of each of the three resulting images are then combined into a single image, that more nearly represents the high dynamic range the eye perceives when viewing the scene.

The problem for classic 8-bit “true” color coding is not merely that 8-bit channels are inadequate for capturing HDR images, but high dynamic range images cannot be reproduced on displays using just 24 bits/pixel. The recent (August 2015) Consumer Electronics Association HDR10 Media Profile standard, for HDR compatible displays, uses 10-bit channels (as the “HDR10” name suggests). Of course, 10-bit channels require displays to process 30 bits of color information for pixels. Which is to say, to make use of their capabilities, content providers must transmit 30-bits/pixel of color information to HDR displays.

Appendix 2: Adaptive Bitrate Streaming

Streaming companies, to make sure streamed content gets through to their customers, take advantage of something known as Adaptive Bitrate Streaming (ABS). ABS leverages the fact that, unlike traditional broadcast technologies, where data flows only one way, from a transmitter to a receiver, the Internet is bidirectional. An Internet provider can not only send data to a customer, but can receive data back. Two-way communication opens up a whole new realm of interactive possibilities, including ABS.

The bidirectional communication involved in ABS consists of sending data to a target address and, in real time, receiving back information about available bandwidth at the target destination. Thus, on getting a video request, the usual protocol is to begin sending a relatively low-bitrate stream to the target address. For this reason, when you first start up an on-demand video, the picture is often not very good.

As video delivery begins, however, the receiving device reports back to the sender on its available bandwidth, allowing the sender to adjust its transmission. Thus, if more bandwidth is available, the sender switches to a higher bitrate feed, improving the initial softer, lower-bitrate video. Or, if even the initial, lower bandwidth feed exceeds the capacity available, the feed starts at a still lower bitrate. Should available bandwidth be below the lowest usable level for the target display, the sender instead puts out a “try again later” message.

A standard ABS ladder, with full HD as its top rung, might consist of 10 steps, ranging, say, from a high resolution of 1920 x 1080 at 5800Kbps to a low resolution of 320 x 240 at 235Kbps. Using this ladder, a client that fails to meet the high requirement of 5800Kbps might still get a (softer) full HD picture, at a lower bandwidth of 4300Kbps. If the client device fails even that bandwidth test, the image will shift from “full HD” to 720p HD (1280 x 720) at 3000Kbps. A softer version of this format might be available at, say, 2350Kbps. Below 2350Kbps, the picture reduces again, to a 720 x 480 SD format, streamed at 1750Kbps. And so on, all the way down to the bottom rung of the ladder, the small screen 320 x 240 mobile format streamed at 235Kbps.

In short, you may have a “full HD” TV, but if all don’t have at least 2350Kbps of streaming bandwidth available when you’re watching House of Cards, you won’t be seeing it in HD. And if you don’t have at least 4300Kbps available, you will see it in 720p HD rather than a full HD 1920 x 1080 format.

This sort of ABS ladder is the secret behind Netflix’s cheerful promotion of 4K content. Yes, they will make House of Cards available in a 4K UHD format. All it takes is adding more UHD rungs to the top of their ABS ladder. But they do not say you, in your particular circumstances, will receive House of Cards in 4K. In fact, they do not guarantee any streaming customer will ever receive the program in this format. In reality, the version customers receive depends entirely on the sort of receiver they have, and the fluctuating bandwidth available at that receiver while the program is playing.

Thus, if you have a UHD receiver hooked up to a pipe with (say) 18Mbps available, Netflix surely will be happy to send you House of Cards in a 4K 3840 x 2160 format. Of course, this will not be the “full 4K” version of the program (with 10-bit HDR color channels streamed at a rate of 60fps, which would require far more bandwidth). But, at least, it will be 4K resolution. However, if no one has 18Mbps available at 9 PM on a Sunday, then no one watching House of Cards at that time will get the UHD version of the program. Yes, many, perhaps all can still see the program. But it will be shown via a lower resolution, lower bitrate stream, as determined by the particular circumstances of each customer.

To repeat, OTT providers just provide the content—in every format, up to and including UHD—but they are not responsible for which format anyone receives. That is determined by fluctuating circumstances, over which they have no control, across the Internet they never built.